

CONSTRUCTION OF BAYESIAN DEFORMABLE MODELS VIA STOCHASTIC APPROXIMATION ALGORITHM: A CONVERGENCE STUDY

S. ALLASSONNIÈRE*, E. KUHN†, AND A. TROUVÉ‡

Abstract. The problem of the definition and the estimation of generative models based on deformable templates from raw data is of particular importance for modeling non-aligned data affected by various types of geometrical variability. This is especially true in shape modeling in the computer vision community or in probabilistic atlas building in Computational Anatomy. A first coherent statistical framework modeling the geometrical variability as hidden variables was described by Allasonnière, Amit and Trouvé in [2]. The present paper gives a theoretical proof of convergence of effective stochastic approximation expectation strategies to estimate such models and shows the robustness of this approach against noise through numerical experiments in the context of handwritten digit modeling.

Key words. stochastic approximation algorithms, non rigid-deformable templates, shapes statistics, Bayesian modeling, MAP estimation.

AMS subject classifications. 60J22, 62F10, 62F15, 62M40.

1. Introduction. In the field of image analysis, the statistical analysis and modeling of variable objects from a limited set of examples is still a quite challenging and a largely unsolved problem and depends strongly on the use of adequate representations of data. One such representation is the so-called dense deformable template (DDT) framework [4]. Observations are defined as deformations, taken from a family of deformations of moderate “dimensionality”, of a given exemplar or template. Such a representation appears particularly adapted to the emerging field of Computational Anatomy where one aims at building statistical models of the anatomical variability within a given population [12]. However, research on DDT has been mainly focused on the variational point of view, in which DDT is used as an efficient vehicle for a wide range of registration algorithms [7]. The problem of template estimation, viewed as a statistical estimation problem of parameters of generative models of images of deformable objects, has received much less attention.

In this paper, we consider the hierarchical Bayesian framework for dense deformable templates developed by Allasonnière, Amit and Trouvé in [2]. Each image in a given population is assumed to be generated as a noisy and randomly deformed version of a common template drawn from a prior distribution on the set of templates. Individual deformations in their framework are treated as *hidden variables* (or equivalently random effects in the mixed effects setting), whereas the template and the law of the deformations are parameters (or equivalently fixed effects) of interest. Parameter estimation for this model could be performed by Maximum A Posteriori (MAP) for which existence and consistency (as the number of parameters observed images tends to infinity) has been proved (see [2]). This contrasts with earlier work in [11] using a penalized likelihood (PL) or the more recent maximum description length approach in [14] for which consistency cannot be proved because the deformations are considered as *nuisance parameters* to be estimated.

Our contribution in this paper is in defining effective and theoretically proven convergent stochastic algorithms for computing (local) maxima of the posterior on the parameters for Bayesian deformable template models. First, we specify an adapted stochastic approximation expectation minimization algorithm (SAEM algorithm) in this highly demanding framework where the hidden variables are non rigid deformation fields living in finite but high dimensional space (typically hundreds or more dimensions). In particular, special attention is needed to the sampling of the posterior distribution on the deformations. Obviously, MCMC samplers are unavoidable, but non adaptive proposal distributions yielding simple symmetric random steps are of limited practical interest. The present paper introduces a more sophisticated hybrid Gibbs sampling

*CMAP Ecole Polytechnique, Route de Saclay, F-91128 Palaiseau, France

†LAGA, Université Paris 13, 99, Av. Jean-Baptiste Clément, F-93430 Villetaneuse, France

‡CMLA, ENS Cachan, CNRS, PRES UniverSud, 61 Av. Président Wilson, F-94230 Cachan, France

scheme allowing an acceptable rejection rate during the Estimation step. The overall algorithm is cast in the larger class of SAEM-MCMC algorithms introduced in [13]. Second, we extend the convergence theory of SAEM-MCMC algorithms developed in [13] to cover the case of *unbounded* random effects arising naturally for deformation fields. The core material of this extension is based on the general stability and convergence results for stochastic algorithms with truncation on random boundaries given in [6]. The main technical point is that in the presence of unbounded random effects and sequential estimation of the covariance matrix of the random effects, the usual regularity conditions of the solutions of the Poisson equations for the Markovian dynamic as a function of the parameters cannot be verified and have to be relaxed. As a result we provide a new general stochastic approximation convergence theorem with a weaker set of assumptions. Third, we prove that the conditions for stability and convergence are fulfilled for our general SAEM-MCMC estimation algorithm of Bayesian dense deformable templates. Indeed, a well known weakness of general stochastic approximation algorithm convergence results is that they rarely provide proofs of convergence for the algorithms *used in practice* since in these implementations the assumptions are not satisfied or hard to verify (see [6]). Since stochastic approximation algorithms have started recently to attract interest for deformable model estimation (see [1] and [16]) our results provide the missing theoretical foundations and guidelines for their effective use. As an illustration of the potential of such SAEM-MCMC approaches in the context of deformable templates, in particular in the presence of noisy data, we present a set of experiments with images of handwritten digits.

This article is organized as follows. Section 2 briefly reviews the hierarchical Bayesian deformable template model proposed by Allasonnière, Amit and Trouvé in [2]. In Section 3, we develop the SAEM-MCMC strategy for the estimation of the parameters. Then in Section 4, we state our general convergence result for truncated stochastic approximation algorithm extending the Andrieu et al. Theorem of convergence in [6] and state that the designed family of SAEM-MCMC algorithms in the previous section satisfy the assumptions. The proof of this last statement is postponed to Section 6 after Section 5 concentrates on experiments. In a final Section, we provide a short discussion and conclusion.

2. Observation model. Let us recall the model introduced in [2]. We are given gray level images $(y_i)_{1 \leq i \leq n}$ observed on a grid of pixels $\{v_u \in D \subset \mathbb{R}^2, u \in \Lambda\}$ which is embedded in a continuous domain $D \subset \mathbb{R}^2$, (typically $D = [-1, 1] \times [-1, 1]$). Although the images are observed only at the pixels $(v_u)_u$, we are looking for a template image $I_0 : \mathbb{R}^2 \rightarrow \mathbb{R}$ defined on the plane (the extension to images on \mathbb{R}^d is straightforward). Each observation y is assumed to be the discretization on a fixed pixel grid of a deformation of the template plus independent noise. Specifically for each observation there exists an *unobserved* deformation field $z : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ such that for $u \in \Lambda$

$$y(u) = I_0(v_u - z(v_u)) + \epsilon(u) ,$$

where ϵ denotes an independent additive noise.

2.1. Models for template and deformation. Our model takes into account two complementary sides: photometry -indexed by p , and geometry -indexed by g . Estimating the template and the distribution on deformations directly as a continuous function would be an infinite dimensional problem. We reduce this problem to a finite dimensional one by restricting the search to a parameterized space of functions. The template $I_0 : \mathbb{R}^2 \rightarrow \mathbb{R}$ and the deformation $z : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ are assumed to belong to fixed reproducing kernel Hilbert spaces V_p and V_g defined by their respective kernels K_p and K_g . Moreover, we restrict them to the subset of linear combinations of the kernels centered at some fixed control points in the domain D : $(v_{p,j})_{1 \leq j \leq k_p}$ respectively $(v_{g,j})_{1 \leq j \leq k_g}$. They are therefore parameterized by the coefficients $\alpha \in \mathbb{R}^{k_p}$ and $\beta \in \mathbb{R}^{k_g} \times \mathbb{R}^{k_g}$ as follows. For all v in D , let

$$I_\alpha(v) \triangleq (\mathbf{K}_p \alpha)(v) \triangleq \sum_{j=1}^{k_p} K_p(v, v_{p,j}) \alpha^j ,$$

and

$$z_\beta(v) \triangleq (\mathbf{K}_g \beta)(v) \triangleq \sum_{j=1}^{k_g} K_g(v, v_{g,j}) \beta^j.$$

Other forms of smooth parametric representations of the images and of the deformation fields could be used without changing the overall results.

2.2. Parametric model. For clarity, we denote by $\mathbf{y}^t = (y_1^t, \dots, y_n^t)$ and by $\beta^t = (\beta_1^t, \dots, \beta_n^t)$ the collection of data and their corresponding deformation coefficients. The statistical model of the observations we consider is a generative hierarchical one. We assume conditional normal distributions for \mathbf{y} and β :

$$\begin{cases} \beta \sim \otimes_{i=1}^n \mathcal{N}_{2k_g}(0, \Gamma_g) \mid \Gamma_g, \\ \mathbf{y} \sim \otimes_{i=1}^n \mathcal{N}_{|\Lambda|}(z_{\beta_i} I_\alpha, \sigma^2 \text{Id}) \mid \beta, \alpha, \sigma^2, \end{cases} \quad (2.1)$$

where \otimes denotes the product of distributions of independent variables and $zI_\alpha(u) = I_\alpha(v_u - z(v_u))$, for u in Λ denotes the action of the deformation on the template image. The parameters of interest are α which determines the template image, σ^2 the variance of the additive noise and Γ_g the covariance matrix of the variables β . We assume that $\theta = (\alpha, \sigma^2, \Gamma_g)$ belongs to an open parameter space Θ :

$$\Theta \triangleq \{ \theta = (\alpha, \sigma^2, \Gamma_g) \mid \alpha \in \mathbb{R}^{k_p}, \|\alpha\| < R, \sigma > 0, \Gamma_g \in \text{Sym}_{2k_g}^+ \},$$

where $\|\cdot\|$ is the Euclidean norm, $\text{Sym}_{2k_g}^+$ is the cone of real positive $2k_g \times 2k_g$ definite symmetric matrices and R an arbitrary positive constant.

The likelihood of the observed data q_{obs} can be written as an integral over the unobserved deformation variables. Let us denote by q_c the conditional likelihood of the observations given the hidden variables and by q_m the likelihood of these missing variables. Then,

$$q_{obs}(\mathbf{y}|\theta) = \int q_c(\mathbf{y}|\beta, \alpha, \sigma^2) q_m(\beta|\Gamma_g) d\beta,$$

where all the densities are determined by the model (2.1).

2.3. Bayesian model. Even though the parameters are finite dimensional, the maximum-likelihood estimator can yield degenerate estimates when the training sample is small. By introducing prior distributions on the parameters, estimation with small samples is still possible. The regularizing effect of such priors can be seen in the parameter update steps (cf. [2]). We use a generative model based on standard conjugate prior distributions for parameters $\theta = (\alpha, \sigma^2, \Gamma_g)$ with fixed hyper-parameters. Specifically, we assume a normal prior for α , an inverse-Wishart prior on σ^2 and an inverse-Wishart prior on Γ_g . Furthermore, all priors are assumed to be independent. This yields $\theta = (\alpha, \sigma^2, \Gamma_g) \sim q_{para} \triangleq \nu_p \otimes \nu_g$ where

$$\begin{cases} \nu_p(d\alpha, d\sigma^2) \propto \exp\left(-\frac{1}{2}(\alpha - \mu_p)^t (\Sigma_p)^{-1} (\alpha - \mu_p)\right) \left(\exp\left(-\frac{\sigma_0^2}{2\sigma^2}\right) \frac{1}{\sqrt{\sigma^2}}\right)^{a_p} d\sigma^2 d\alpha, \quad a_p \geq 3, \\ \nu_g(d\Gamma_g) \propto \left(\exp(-\langle \Gamma_g^{-1}, \Sigma_g \rangle_F / 2) \frac{1}{\sqrt{|\Gamma_g|}}\right)^{a_g} d\Gamma_g, \quad a_g \geq 4k_g + 1. \end{cases} \quad (2.2)$$

For two matrices A and B , we define $\langle A, B \rangle_F \triangleq \text{tr}(A^t B)$ the Frobenius dot product on the set of matrices where tr denotes the trace of the matrix.

3. Parameter estimation based on stochastic approximation EM. In our Bayesian framework, we obtain from [2] the existence of the MAP estimator

$$\tilde{\theta}_n = \underset{\theta \in \Theta}{\text{argmax}} q_B(\theta|\mathbf{y}),$$

where q_B denotes the posterior likelihood of the parameters given the observations. The dependence on n refers to the sample size.

We turn now to the maximization problem of the penalized posterior distribution $q_B(\theta|\mathbf{y})$ which has no closed form in our case. Indeed, the probability density function is known up to a renormalization constant. That prevents a direct computation of $\tilde{\theta}_n$.

In order to solve this problem, we apply an “EM like” algorithm to approximate the MAP estimator $\tilde{\theta}_n$. The solution we propose is to base our algorithm on the use of the Stochastic Approximation EM (SAEM). First, we outline certain characteristics of our model, which highlight the reasons for the choice of the particular procedure and enable us to simplify its implementation.

3.1. Model characteristics. An important characteristic of our model is that it belongs to the curved exponential family. In other words the complete likelihood q can be written as:

$$q(\mathbf{y}, \boldsymbol{\beta}, \theta) = \exp[-\psi(\theta) + \langle S(\boldsymbol{\beta}), \phi(\theta) \rangle],$$

where the sufficient statistic S is a Borel function on \mathbb{R}^N , with $N \triangleq 2nk_g$, taking its values in an open subset \mathcal{S} of \mathbb{R}^m and ψ, ϕ two Borel functions on Θ . (Note that S, ϕ and ψ may depend also on \mathbf{y} , but since \mathbf{y} will stay fixed in what follows, we omit this dependence).

In our setting, we obtain the following formula:

$$\log q(\mathbf{y}, \boldsymbol{\beta}, \theta) = \log q_c(\mathbf{y}|\boldsymbol{\beta}, \theta) + \log q_m(\boldsymbol{\beta}|\theta) + \log q_{para}(\theta),$$

where q_{para} denotes the prior density of the parameters defined in the previous paragraph.

For any $1 \leq j \leq k_p$ and any $u \in \Lambda$, we denote by

$$K_p^\beta(u, j) = K_p(v_u - z_\beta(v_u), v_{p,j})$$

the matrix which corresponds to the deformation of the kernel K_p through z_β at pixel u and evaluated at pixel location

locpixel_u. Then, for some constant C independent of θ ,

$$\begin{aligned} \log q(\mathbf{y}, \boldsymbol{\beta}, \theta) &= \sum_{i=1}^n \left\{ -\frac{|\Lambda|}{2} \log(\sigma^2) - \frac{1}{2\sigma^2} \|y_i - K_p^{\beta_i} \alpha\|^2 \right\} \\ &+ \sum_{i=1}^n \left\{ -\frac{1}{2} \log(|\Gamma_g|) - \frac{1}{2} \beta_i^t \Gamma_g^{-1} \beta_i \right\} \\ &+ a_g \left\{ -\frac{1}{2} \log(|\Gamma_g|) - \frac{1}{2} \langle \Gamma_g^{-1}, \Sigma_g \rangle_F \right\} - \frac{1}{2} (\alpha - \mu_p)^t \Sigma_p^{-1} (\alpha - \mu_p) \\ &+ a_p \left\{ -\frac{1}{2} \log(\sigma^2) - \frac{\sigma_0^2}{2\sigma^2} \right\} + C. \end{aligned}$$

Note that $\|y_i - K_p^{\beta_i} \alpha\|^2 = (y_i - K_p^{\beta_i} \alpha)^t (y_i - K_p^{\beta_i} \alpha)$, where $K_p^{\beta_i} \alpha$ is another way to write the action of the deformation z_{β_i} on the template I_α denoted previously by $z_{\beta_i} I_\alpha$. This form emphasizes the dot product between the sufficient statistics and a function of the parameters. It can be easily verified that the following matrix-valued functions are the sufficient statistics (up to a multiplicative constant) :

$$\begin{aligned} S_1(\boldsymbol{\beta}) &= \sum_{1 \leq i \leq n} (K_p^{\beta_i})^t y_i, \\ S_2(\boldsymbol{\beta}) &= \sum_{1 \leq i \leq n} (K_p^{\beta_i})^t (K_p^{\beta_i}), \\ S_3(\boldsymbol{\beta}) &= \sum_{1 \leq i \leq n} \beta_i^t \beta_i. \end{aligned}$$

For simplicity, we denote $S(\beta) = (S_1(\beta), S_2(\beta), S_3(\beta))$ for any $\beta \in \mathbb{R}^N$ and define the sufficient statistic space as

$$\mathcal{S} = \left\{ (S_1, S_2, S_3) \mid S_1 \in \mathbb{R}^{k_p}, S_2 + \sigma_0^2 \Sigma_p^{-1} \in \text{Sym}_{k_p}^+, S_3 + a_g \Sigma_g \in \text{Sym}_{2k_g}^+ \right\}.$$

Identifying S_2 and S_3 with their lower triangular parts, the set \mathcal{S} can be viewed as an open set of \mathbb{R}^{n_s} with $n_s = k_p + \frac{k_p(k_p+1)}{2} + k_g(2k_g + 1)$.

In [2], the existence of the parameter estimate $\hat{\theta}(S)$ that maximizes the complete log-likelihood has been proved. It can easily be shown that α , σ^2 and Γ_g are explicitly expressed with the above sufficient statistics as follows:

$$\begin{cases} \Gamma_g(S) &= \frac{1}{n+a_g}(S_3 + a_g \Sigma_g), \\ \alpha(S) &= (S_2 + \sigma^2(S)(\Sigma_p)^{-1})^{-1} (S_1 + \sigma^2(S)(\Sigma_p)^{-1} \mu_p), \\ \sigma^2(S) &= \frac{1}{n|\Lambda|+a_p} (n\|\mathbf{y}\|^2 + \alpha(S)^t S_2 \alpha(S) - 2\alpha(S)^t S_1 + a_p \sigma_0^2). \end{cases} \quad (3.1)$$

All these formulas also prove the smoothness of $\hat{\theta}$ on the subset \mathcal{S} .

3.2. SAEM-MCMC algorithm with truncation on random boundaries. In order to compute the MAP estimator for our Bayesian model, we use a variant of the EM (Expectation-Maximization, [9]) algorithm. This algorithm is quite natural when we have to maximize a likelihood under a hierarchical model with missing variables. Unfortunately, direct computation is not tractable and we have to find a solution to overcome the problematic E step where we have to compute an expectation with respect to the posterior distribution on β given \mathbf{y} . A first attempt was proposed in [2] where this conditional distribution is approximated by a Dirac distribution at its mode (Fast Approximation with Mode -FAM-EM). The results are very interesting, however, the authors point out the lack of convergence of the FAM-EM algorithm when the quality of the input images is not good, typically when they are noisy. This is the issue we consider here. We propose an algorithm that ensures the convergence of the resulting sequence of estimators toward the MAP whatever the quality of the input.

This solution is a procedure combining the Stochastic Approximation EM (SAEM) with Markov Chain Monte Carlo (MCMC) in a more general framework than that proposed by [13], which in turn generalized the algorithm introduced by [8]. Indeed, the k^{th} iteration of the SAEM-MCMC algorithm consists of three steps:

Step 1 : Simulation step. The missing data, i.e. the deformation parameters β , are drawn using the transition probability of a convergent Markov chain Π_θ having the posterior distribution $q_{post}(\cdot|\mathbf{y}, \theta)$ as its stationary distribution:

$$\beta_k \sim \Pi_{\theta_{k-1}}(\beta_{k-1}, \cdot).$$

Step 2 : Stochastic approximation step. A stochastic approximation is done on the complete log-likelihood using the simulated value of the missing data:

$$Q_k(\theta) = Q_{k-1}(\theta) + \Delta_{k-1}[\log q(\mathbf{y}, \beta_k, \theta) - Q_{k-1}(\theta)],$$

where $\Delta = (\Delta_k)_k$ is a decreasing sequence of positive step-sizes.

Step 3 : Maximization step. The parameters are updated in the M-step:

$$\theta_k = \underset{\theta \in \Theta}{\operatorname{argmax}} Q_k(\theta).$$

The initial values Q_0 and θ_0 are arbitrarily chosen.

REMARK 1. *We cannot use the direct SAEM algorithm. Indeed, this would require sampling the hidden variable from the posterior distribution which is known only up to a normalization*

constant. This sampling is not possible here due to the complexity of the posterior probability density function.

Since our model belongs to the curved exponential family, the stochastic approximation step can easily be done on the sufficient statistics S instead of on the complete log-likelihood. Then the maximization step 3 is straightforward, replacing in (3.1) the sufficient statistics with their corresponding stochastic approximations.

The convergence of this algorithm has been proved in [13] in the particular case of missing variables living in a compact subset of \mathbb{R}^N . However, as we set a Gaussian prior on the missing variables β , we cannot assume that their support is compact. In order to provide an algorithm whose convergence can be proved in the current framework we have to use a more general setting introduced in [6] which involves truncation on random boundaries. The proof is given in Section 4. This can be formalized as follows.

Let $(\mathcal{K}_q)_{q \geq 0}$ be a sequence of increasing compact subsets of \mathcal{S} such as $\cup_{q \geq 0} \mathcal{K}_q = \mathcal{S}$ and $\mathcal{K}_q \subset \text{int}(\mathcal{K}_{q+1})$, for all $q \geq 0$. Let $\varepsilon = (\varepsilon_k)_{k \geq 0}$ be a monotone non-increasing sequence of positive numbers and K a compact subset of \mathbb{R}^N . We construct a sequence $((\beta_k, s_k))_{k \geq 0}$ as described in Algorithm 1 as follows. As long as the stochastic approximation does not wander outside the current compact set and is not too far from its previous value, we run the SAEM-MCMC algorithm. As soon as one of these conditions is not satisfied, we reinitialize the sequences of β and s using a projection (for more details see [6]), we increase the size of the compact set and continue the iterations until convergence. This is detailed in the following steps :

Initialization step : Initialize β_0 and s_0 in two fixed compact sets K and \mathcal{K}_0 respectively.

Then, for the k^{th} iteration, repeat the following four steps :

Step 1 : MCMC simulation step . Draw one new element $\bar{\beta}$ of the non-homogeneous Markov Chain with respect to the kernel with the current parameters $\Pi_{\theta_{k-1}}$ and starting at β_{k-1} .

$$\bar{\beta} \sim \Pi_{\theta_{k-1}}(\beta_{k-1}, \cdot).$$

Step 2 : Stochastic approximation step . Compute

$$\bar{s} = s_{k-1} + \Delta_{\zeta_{k-1}}(S(\bar{\beta}) - s_{k-1}). \quad (3.2)$$

Step 3 : Truncation on random boundaries . If \bar{s} is outside the current compact set $\mathcal{K}_{\kappa_{k-1}}$ or too far from the previous value s_k , then restart the stochastic approximation in the initial compact set, extend the truncation boundary to \mathcal{K}_{κ_k} and start again with a bounded value of the missing variable. Otherwise, set $(\beta_k, s_k) = (\bar{\beta}, \bar{s})$ and keep the truncation boundary to $\mathcal{K}_{\kappa_{k-1}}$.

Step 4 : Maximization step . Update the parameters using (3.1).

In this algorithm, the MCMC simulation step has to be explained since it involves the choice of the transition kernel of the Markov chain. Usually, one uses a Metropolis-Hastings algorithm in which a candidate value is sampled from a proposal distribution followed by an accept-reject step. However, there are different possible proposal distributions. The only requirement is that all these kernels lead to an ergodic Markov chain whose stationary distribution is our posterior distribution. The choice among these possibilities should be based on the specific framework we are working in.

While minimizing the Kullback-Leibler distance between the stationary distribution $\beta \rightarrow \pi_{\theta}(\beta)$ and a tensorial product $\beta \rightarrow \otimes_{i=1}^n p(\beta_i)$ corresponding to independent identically distributed missing variables, we get that p is proportional to $\frac{1}{n} \sum_{i=1}^n q_{post}(\cdot | y_i, \theta)$. As n tends to ∞ and for a

given θ , p converges a.s. towards the prior pdf on the missing variable $q_m(\cdot|\theta)$. This suggest to use as proposal the prior distribution which involves the current parameters.

On the other hand, the setting we have in this paper deals with high dimensional missing variables. This raises several issues. If we simulate candidates for the hidden variable as a complete vector, it appears that most of the candidates are rejected. This is a typical high dimensional concentration phenomenon : locally around a current point, the proportion of the space occupied by acceptable moves becomes negligible when the space dimension grows. From a more practical point of view, even if the proposed candidate is drawn with respect to the current prior distribution, it creates a deformation that is very different from the current one and too large for the corresponding deformed template to fit the observations. This yields very few possible moves from the current missing variable value and the algorithm is stuck in a non-optimal location or converges very slowly.

One solution is to update the chain one coordinate at a time conditionally on the others. This corresponds to a Gibbs sampler and leads to more relevant candidates which have a higher chance to be accepted (cf. [3]). From an image analysis point of view, this put stronger conditions on the kind of deformations which are produced when proposing a candidate for each coordinate. Knowing the tendency of the movement given by the other coordinates, the candidate will either confirm it or not depending if this is a suitable movement. It will thus be accepted with a corresponding probability. Even if some coordinates remain unchanged, some others are updated which enables the algorithm to visit a larger part of the missing variable support.

Algorithm 1 Stochastic approximation with truncation on random boundaries

```

Set  $\beta_0 \in K$ ,  $s_0 \in \mathcal{K}_0$ ,  $\kappa_0 = 0$ ,  $\zeta_0 = 0$ , and  $\nu_0 = 0$ .
for all  $k \geq 1$  do
  compute  $\bar{s} = s_{k-1} + \Delta_{\zeta_{k-1}}(S(\bar{\beta}) - s_{k-1})$ 
  where  $\bar{\beta}$  is sampled from a transition kernel  $\Pi_{\theta_{k-1}}(\beta_{k-1}, \cdot)$ .
  if  $\bar{s} \in \mathcal{K}_{\kappa_{k-1}}$  and  $\|\bar{s} - s_{k-1}\| \leq \varepsilon_{\zeta_{k-1}}$  then
    set  $(\beta_k, s_k) = (\bar{\beta}, \bar{s})$  and  $\kappa_k = \kappa_{k-1}$ ,  $\nu_k = \nu_{k-1} + 1$ ,  $\zeta_k = \zeta_{k-1} + 1$ 
  else
    set  $(\beta_k, s_k) = (\bar{\beta}, \bar{s}) \in K \times \mathcal{K}_0$  and  $\kappa_k = \kappa_{k-1} + 1$ ,  $\nu_k = 0$ ,  $\zeta_k = \zeta_{k-1} + \phi(\nu_{k-1})$ 
    where  $\phi: \mathbb{N} \rightarrow \mathbb{Z}$  is a function such that  $\phi(k) > -k$  for any  $k$ 
    and  $(\bar{\beta}, \bar{s})$  can be chosen through different ways (cf. [6]).
  end if
   $\theta_k = \hat{\theta}(s_k)$ .
end for
```

REMARK 2. The index κ denotes the current active truncation set, the index ζ is the current index in the sequences Δ and ε and the index ν denotes the number of iterations since the last projection.

3.3. Transition probability of the Markov chain. We now explain how to simulate the missing variables thanks to a Markov Chain Monte Carlo algorithm having the posterior distribution as its stationary distribution. Due to the inherent high dimensionality N of β , we consider a Gibbs sampler to sequentially scan all coordinates β^j for $1 \leq j \leq N$.

Denote by $\beta^{-j} = (\beta^l)_{l \neq j}$. We consider here a hybrid Gibbs sampler i.e. each step of the Gibbs sampler includes a Metropolis-Hastings step. The proposal law is chosen as $q_j(\cdot|\beta^{-j}, \theta)$ i.e. the conditional law based on the current parameter value θ derived from the normal distribution q_m .

If b is a proposed value at coordinate j , the acceptance rate of the Metropolis-Hastings algorithm is given by

$$r_j(\beta^j, b; \beta^{-j}, \theta) = \left[\frac{q_j(b|\beta^{-j}, \mathbf{y}, \theta) q_j(\beta^j|\beta^{-j}, \theta)}{q_j(\beta^j|\beta^{-j}, \mathbf{y}, \theta) q_j(b|\beta^{-j}, \theta)} \wedge 1 \right].$$

Since

$$q_j(\beta^j | \beta^{-j}, \mathbf{y}, \theta) \propto q_{obs}(\mathbf{y} | \beta, \theta) q_j(\beta^j | \beta^{-j}, \theta),$$

the acceptance rate can be simplified to

$$r_j(\beta^j, b; \beta^{-j}, \theta) = \left[\frac{q_{obs}(\mathbf{y} | \beta_{b \rightarrow j}, \theta)}{q_{obs}(\mathbf{y} | \beta, \theta)} \wedge 1 \right],$$

where for any $b \in \mathbb{R}$ and $1 \leq j \leq N$, we denote by $\beta_{b \rightarrow j}$ the unique vector which is equal to β everywhere except at coordinate j where it equals b . An illustration of the hybrid Gibbs sampler can be found in [17]. The following steps are performed for each coordinate j :

Step 1 : Proposition . Sample b with respect to the density $q_j(\cdot | \beta^{-j}, \theta)$.

Step 2 : Accept-reject . Compute $r_j(\beta^j, b; \beta^{-j}, \theta)$ and with probability $r_j(\beta^j, b; \beta^{-j}, \theta)$, update β^j to b .

In Algorithm 2, we summarize the transition step of the Markov chain.

Algorithm 2 Transition step $k \rightarrow k+1$ using a hybrid Gibbs sampler

Require: $\beta = \beta_k; \theta = \theta_k$

Gibbs sampler:

for all $j = 1 : N$ do

Metropolis-Hastings procedure:

$b \sim q_j(\cdot | \beta^{-j}, \theta)$;

Compute $r_j(\beta^j, b; \beta^{-j}, \theta) = \left[\frac{q_{obs}(\mathbf{y} | \beta_{b \rightarrow j}, \theta)}{q_{obs}(\mathbf{y} | \beta, \theta)} \wedge 1 \right]$

With probability $r_j(\beta^j, b; \beta^{-j}, \theta)$, update β^j : $\beta^j \leftarrow b$

end for

This yields the transition probability kernel of our Markov chain on β : for coordinate j , the kernel is

$$\begin{aligned} \Pi_{\theta,j}(\beta, d\mathbf{z}) = & (\otimes_{m \neq j} \delta_{\beta^m}(d\mathbf{z}^m)) \times \left[q_j(d\mathbf{z}^j | \beta^{-j}, \theta) r_j(\beta^j, d\mathbf{z}^j; \beta^{-j}, \theta) + \right. \\ & \left. \delta_{\beta^j}(d\mathbf{z}^j) \int (1 - r_j(\beta^j, b; \beta^{-j}, \theta)) q_j(b | \beta^{-j}, \theta) db \right] \quad (3.3) \end{aligned}$$

and $\Pi_\theta = \Pi_{\theta,N} \circ \dots \circ \Pi_{\theta,1}$ is therefore the kernel associated with a complete scan.

4. Convergence analysis. We prove a general theorem on the convergence of stochastic approximations for which our algorithm convergence is a special case.

The hybrid Gibbs sampler used to generate the ergodic Markov chain does not satisfy some of the assumptions of the convergence result presented in [6]. We therefore weaken some of their conditions, introducing an absorbing set for the stochastic approximation and weakening their Hölder conditions on some functions of the Markov chain.

4.1. Stochastic approximation convergence Theorem. Let \mathcal{S} be a subset of \mathbb{R}^{n_s} for some integer n_s . Let X be a measurable space. For all $s \in \mathcal{S}$ let $H_s : X \rightarrow \mathcal{S}$ be a measurable function. Let $\Delta = (\Delta_k)_k$ be a sequence of positive step-sizes.

Define the stochastic approximation sequence $(s_k)_k$ as follows :

$$\begin{cases} s_k = s_{k-1} + \Delta_{k-1} H_{s_{k-1}}(\beta_k) & \text{with } \beta_k \sim \Pi_{s_{k-1}}(\beta_{k-1}, \cdot), \quad \text{if } s_{k-1} \in \mathcal{S} \\ s_k = s_c & \text{with } \beta_k = \beta_c, \quad \text{if } s_{k-1} \notin \mathcal{S}, \end{cases} \quad (4.1)$$

where $s_c \notin \mathcal{S}$, $\beta_c \notin X$ and $(\Pi_s)_{s \in \mathcal{S}}$ is a family of Markov transition probabilities on X . Denote by Q_Δ the transition which generates $((\beta_k, s_k))_k$. We consider the natural filtration of the *non-homogeneous* chain $((\beta_k, s_k))_k$ and denote respectively by $\mathbb{P}_{\beta,s}^\Delta$ and $\mathbb{E}_{\beta,s}^\Delta$ the probability measure

and the corresponding expectation generated by this Markov chain starting at (β, s) and using the sequence Δ .

If the transition kernel Π_s of the Markov chain admits a stationary distribution π_s and if for any $s \in \mathcal{S}$, H_s is integrable with respect to π_s , then we denote by h the mean field associated with our stochastic approximation so that :

$$h(s) = \int H_s(\beta) \pi_s(\beta) d\beta.$$

The algorithm defined in 4.1 is usually designed to solve the equation $h(s) = 0$ where h is called the mean field function.

Let $(\mathcal{K}_q)_{q \geq 0}$ be a sequence of increasing compact subsets of \mathcal{S} such as $\cup_{q \geq 0} \mathcal{K}_q = \mathcal{S}$ and $\mathcal{K}_q \subset \text{int}(\mathcal{K}_{q+1}), \forall q \geq 0$. Let $\varepsilon = (\varepsilon_k)_{k \geq 0}$ be a monotone non-increasing sequence of positive numbers and K a subset of X .

Let $\Phi : X \times \mathcal{S} \rightarrow K \times \mathcal{K}_0$ be a measurable function and $\phi : \mathbb{N} \rightarrow \mathbb{Z}$ be a function such that $\phi(k) > -k$ for any k . Define the *homogeneous* Markov chain

$$(Z_k = (\beta_k, s_k, \kappa_k, \zeta_k, \nu_k))_k \quad (4.2)$$

on $\mathcal{Z} \triangleq X \times \mathcal{S} \times \mathbb{N}^3$ with the following transition at iteration k :

- If $\nu_{k-1} = 0$ then draw $(\beta_k, s_k) \sim Q_{\Delta_{\zeta_{k-1}}}(\Phi(\beta_{k-1}, s_{k-1}), \cdot)$; otherwise draw $(\beta_k, s_k) \sim Q_{\Delta_{\zeta_{k-1}}}((\beta_{k-1}, s_{k-1}), \cdot)$;
- If $\|s_k - s_{k-1}\| \leq \varepsilon_{\zeta_{k-1}}$ and $s_k \in \mathcal{K}_{\kappa_{k-1}}$ then set $\kappa_k = \kappa_{k-1}$, $\zeta_k = \zeta_{k-1} + 1$ and $\nu_k = \nu_{k-1} + 1$; otherwise set $\kappa_k = \kappa_{k-1} + 1$, $\zeta_k = \zeta_{k-1} + \phi(\nu_{k-1})$ and $\nu_k = 0$.

Consider the following assumptions, generalized from [6]. Define for any $V : X \rightarrow [1, \infty]$ and any $g : X \rightarrow \mathbb{R}^{n_s}$ the norm

$$\|g\|_V = \sup_{\beta \in X} \frac{\|g(\beta)\|}{V(\beta)}.$$

A1'. \mathcal{S} is an open subset of \mathbb{R}^{n_s} , $h : \mathcal{S} \rightarrow \mathbb{R}^{n_s}$ is continuous and there exists a continuously differentiable function $w : \mathcal{S} \rightarrow [0, \infty[$ with the following properties.

- (i) There exists an $M_0 > 0$ such that

$$\mathcal{L} \triangleq \{s \in \mathcal{S}, \langle \nabla w(s), h(s) \rangle = 0\} \subset \{s \in \mathcal{S}, w(s) < M_0\}.$$

- (ii) There exists a closed convex set $\mathcal{S}_a \subset \mathcal{S}$ for which $s \rightarrow s + \rho H_s(\beta) \in \mathcal{S}_a$ for any $\rho \in [0, 1]$ and $(\beta, s) \in X \times \mathcal{S}_a$ (\mathcal{S}_a is absorbing) and such that for any $M_1 \in]M_0, \infty]$, the set $\mathcal{W}_{M_1} \cap \mathcal{S}_a$ is a compact set of \mathcal{S} where $\mathcal{W}_{M_1} \triangleq \{s \in \mathcal{S}, w(s) \leq M_1\}$.

- (iii) For any $s \in \mathcal{S} \setminus \mathcal{L}$ $\langle \nabla w(s), h(s) \rangle < 0$.

- (iv) The closure of $w(\mathcal{L})$ has an empty interior.

A2. For any $s \in \mathcal{S}$, the Markov kernel Π_s has a single stationary distribution π_s , $\pi_s \Pi_s = \pi_s$. In addition for all $s \in \mathcal{S}$, $H_s : X \rightarrow \mathcal{S}$ is measurable and $\int_X \|H_s(\beta)\| \pi_s(d\beta) < \infty$.

A3'. For any $s \in \mathcal{S}$, the Poisson equation $g - \Pi_s g = H_s - \pi_s(H_s)$ has a solution g_s . There exist a function $V : X \rightarrow [1, \infty]$ such that $\{\beta \in X, V(\beta) < \infty\} \neq \emptyset$, constants $a \in]0, 1]$, $q \geq 1$ and $p \geq 2$ such that for any compact subset $\mathcal{K} \subset \mathcal{S}$,

- (i)

$$\sup_{s \in \mathcal{K}} \|H_s\|_V < \infty, \quad (4.3)$$

$$\sup_{s \in \mathcal{K}} (\|g_s\|_V + \|\Pi_s g_s\|_V) < \infty, \quad (4.4)$$

- (ii)

$$\sup_{s, s' \in \mathcal{K}} \|s - s'\|^{-a} \{\|g_s - g_{s'}\|_{V^q} + \|\Pi_s g_s - \Pi_{s'} g_{s'}\|_{V^q}\} < \infty. \quad (4.5)$$

- (iii) Let k_0 be an integer. There exist an $\bar{\varepsilon} > 0$ and a constant C such that for any sequence $\varepsilon = (\varepsilon_k)_{k \geq 0}$ satisfying $0 < \varepsilon_k \leq \bar{\varepsilon}$ for all $k \geq k_0$, for any sequence $\Delta = (\Delta_k)_{k \geq 0}$ and for any $\beta \in X$,

$$\sup_{s \in \mathcal{K}} \sup_{k \geq 0} \mathbb{E}_{\beta, s}^{\Delta} [V^{pq}(\beta_k) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k}] \leq CV^{pq}(\beta), \quad (4.6)$$

where $\nu(\varepsilon) = \inf\{k \geq 1, \|s_k - s_{k-1}\| \geq \varepsilon_k\}$ and $\sigma(\mathcal{K}) = \inf\{k \geq 1, s_k \notin \mathcal{K}\}$ and the expectation is related to the non-homogeneous Markov chain $((\beta_k, s_k))_{k \geq 0}$ using the step-size sequence $(\Delta_k)_{k \geq 0}$.

- A4.** The sequences $\Delta = (\Delta_k)_{k \geq 0}$ and $\varepsilon = (\varepsilon_k)_{k \geq 0}$ are non-increasing, positive and satisfy:
 $\sum_{k=0}^{\infty} \Delta_k = \infty$, $\lim_{k \rightarrow \infty} \varepsilon_k = 0$ and $\sum_{k=1}^{\infty} \{\Delta_k^2 + \Delta_k \varepsilon_k^a + (\Delta_k \varepsilon_k^{-1})^p\} < \infty$, where a and p are defined in **(A3')**.

THEOREM 4.1 (General Convergence Result for Truncated Stochastic Approximation). *Assume **(A1')**, **(A2)**, **(A3')** and **(A4)**. Let $\mathcal{K} \subset X$ be such that $\sup_{\beta \in \mathcal{K}} V(\beta) < \infty$ and $\mathcal{K}_0 \subset \mathcal{W}_{M_0} \cap \mathcal{S}_a$ (where M_0 is defined in **(A1')**), and let $(Z_k)_{k \geq 0}$ be the sequence defined in equation (4.2). Then, for all $\beta_0 \in \mathcal{K}$ and $s_0 \in \mathcal{K}_0$, we have $\lim_{k \rightarrow \infty} d(s_k, \mathcal{L}) = 0$ $\mathbb{P}_{\beta_0, s_0, 0, 0, 0}$ -a.s, where $\mathbb{P}_{\beta_0, s_0, 0, 0, 0}$ is the probability measure associated with the chain $(Z_k = (\beta_k, s_k, \kappa_k, \zeta_k, \nu_k))_{k \geq 0}$ starting at $(\beta_0, s_0, 0, 0, 0)$.*

Proof. • The deterministic results obtained by [6] under their assumption **(A1)** remain true if we suppose the existence of an absorbing set as defined in assumption **(A1')**. Indeed, the proofs in [6] can be carried through in the same way restricting the sequences to the absorbing set. Therefore we obtain the same properties. The first one (stated in Lemma 2.1 of [6]) gives the contraction property of the Lyapunov function w . Then, we have (as in Theorem 2.2 of [6]) the fact that a sequence of stochastic approximations stays almost surely in a compact set under some conditions on the perturbation. Lastly, we establish the convergence of such a stochastic approximation.

• We then state a relation between the homogeneous and non-homogeneous chains as done in Lemma 4.1 of [6].

• We now prove an equivalent version of Proposition 5.2 of [6] under our conditions. Indeed the upper bound on the fluctuations of the noise sequence stated in this proposition is relaxed in our case, involving a different power on the function V .

PROPOSITION 4.2. *Assume **(A3')**. Let \mathcal{K} be a compact subset of \mathcal{S} and let $\Delta = (\Delta_k)_k$ and $\varepsilon = (\varepsilon_k)_k$ be two non-increasing sequences of positive numbers such that $\lim_{k \rightarrow \infty} \varepsilon_k = 0$. Then, for p defined in **(A3')**,*

1. *there exists a constant C such that, for any $(\beta, s) \in X \times \mathcal{K}$, any integer l , any $\delta > 0$*

$$\mathbb{P}_{\beta, s}^{\Delta} \left(\sup_{n \geq l} \|S_{l, n}(\varepsilon, \Delta, \mathcal{K})\| \geq \delta \right) \leq C\delta^{-p} \left\{ \left(\sum_{k=l}^{\infty} \Delta_k^2 \right)^{p/2} + \left(\sum_{k=l}^{\infty} \Delta_k \varepsilon_k^a \right)^p \right\} V^{pq}(\beta),$$

where $S_{l, n}(\varepsilon, \Delta, \mathcal{K}) \triangleq \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq n} \sum_{k=l}^n \Delta_k (H_{s_{k-1}}(\beta_k) - h(s_{k-1}))$ and $\mathbb{P}_{\beta, s}^{\Delta}$ is the probability measure generated by the non homogeneous Markov chain $((\beta_k, s_k))_k$ started from the initial condition (β, s) ;

2. *there exists a constant C such that for any $(\beta, s) \in X \times \mathcal{K}$*

$$\mathbb{P}_{\beta, s}^{\Delta} (\nu(\varepsilon) < \sigma(\mathcal{K})) \leq C \left\{ \sum_{k=l}^{\infty} (\Delta_k \varepsilon_k^{-1})^p \right\} V^{pq}(\beta).$$

Proof. The proof of this proposition can proceed as in [6] except for the upper bound on the term involving the Hölder property (second term in the following). Under **(A3'ii)**, this upper bound brings into play an exponent pq on the function V .

Indeed, rewrite $S_{1,n}(\varepsilon, \Delta, \mathcal{K})$ using the Poisson equation and decompose it into a sum of the following five terms :

$$T_n^{(1)} = \sum_{k=1}^n \Delta_k (g_{s_{k-1}}(\beta_k) - \Pi_{s_{k-1}} g_{s_{k-1}}(\beta_{k-1})) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k\}} \quad (4.7)$$

$$T_n^{(2)} = \sum_{k=1}^{n-1} \Delta_{k+1} (\Pi_{s_k} g_{s_k}(\beta_k) - \Pi_{s_{k-1}} g_{s_{k-1}}(\beta_k)) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k+1\}} \quad (4.8)$$

$$T_n^{(3)} = \sum_{k=1}^{n-1} (\Delta_{k+1} - \Delta_k) \Pi_{s_{k-1}} g_{s_{k-1}}(\beta_k) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k+1\}} \quad (4.9)$$

$$T_n^{(4)} = \Delta_1 \Pi_{s_0} g_{s_0}(\beta_0) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq 1\}} - \Delta_n \Pi_{s_{n-1}} g_{s_{n-1}}(\beta_n) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq n\}} \quad (4.10)$$

$$T_n^{(5)} = - \sum_{k=1}^{n-1} \Delta_k \Pi_{s_{k-1}} g_{s_{k-1}}(\beta_k) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) = k\}} . \quad (4.11)$$

We evaluate bounds for the first four quantities. Using the Minkowski inequality for $p/2 \geq 1$ and the Burkholder inequality (for $T_n^{(1)}$) we have :

$$\sup_{s \in \mathcal{S}} \mathbb{E}_{\beta_0, s}^{\Delta} \left[\sup_{n \geq 0} \|T_n^{(1)}\|^p \right] \leq C \left(\sum_{k=1}^{\infty} \Delta_k^2 \right)^{p/2} \sup_{s \in \mathcal{S}} \sum_k \mathbb{E}_{\beta_0, s}^{\Delta} [V^p(\beta_k) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k\}}] , \quad (4.12)$$

$$\sup_{s \in \mathcal{S}} \mathbb{E}_{\beta_0, s}^{\Delta} \left[\sup_{n \geq 0} \|T_n^{(2)}\|^p \right] \leq C \left(\sum_{k=1}^{\infty} \Delta_k \varepsilon_k^{\alpha} \right)^p \sup_{s \in \mathcal{S}} \sum_k \mathbb{E}_{\beta_0, s}^{\Delta} [V^{pq}(\beta_k) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k\}}] , \quad (4.13)$$

$$\sup_{s \in \mathcal{S}} \mathbb{E}_{\beta_0, s}^{\Delta} \left[\sup_{n \geq 0} \|T_n^{(3)}\|^p \right] \leq C \Delta_1^p \sup_{s \in \mathcal{S}} \sum_k \mathbb{E}_{\beta_0, s}^{\Delta} [V^p(\beta_k) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k\}}] , \quad (4.14)$$

$$\sup_{s \in \mathcal{S}} \mathbb{E}_{\beta_0, s}^{\Delta} \left[\sup_{n \geq 0} \|T_n^{(4)}\|^p \right] \leq C \left(\sum_{k=1}^{\infty} \Delta_k^2 \right)^{p/2} \sup_{s \in \mathcal{S}} \sum_k \mathbb{E}_{\beta_0, s}^{\Delta} [V^p(\beta_k) \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k\}}] . \quad (4.15)$$

where C is a constant which depends only upon the compact set \mathcal{K} . The higher power pq appears because of the Hölder condition we assume on the solution of the Poisson equation.

Since now $T_n^{(5)} \mathbb{1}_{\{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq n\}} = 0$ and noting that $V(\beta) \geq 1$, $\forall \beta \in X$, we have $V(\beta)^p \leq V^{pq}(\beta)$. Applying successively (as in [6]) the Markov inequality, condition (4.6) and the Markov property to these upper bounds concludes the proof of the first part of Proposition 4.2.

Concerning the second part, it follows from the same trick as above for upper-bounding the expectation of V^p by V^{pq} .

This ends the proof of the proposition. \square

It is now straightforward to prove the following proposition which corresponds to Proposition 5.3 in [6].

PROPOSITION 4.3. *Assume **(A3')** and **(A4)**. Then, for any subset $K \subset X$ such that $\sup_{\beta \in K} V(\beta) < \infty$, any $M \in (M_0, M_1]$ and any $\delta > 0$, we have $\lim_{k \rightarrow \infty} A(\delta, \varepsilon^{\leftarrow k}, M, \Delta^{\leftarrow k}) = 0$ where $\varepsilon^{\leftarrow k}$ stands for the sequence ε delayed by k switches ($\varepsilon_l^{\leftarrow k} = \varepsilon_{k+l}$ for all $l \in \mathbb{N}$) and*

$$A(\delta, \varepsilon, M, \Delta) = \sup_{s \in \mathcal{K}_0} \sup_{\beta \in K} \left\{ \mathbb{P}_{\beta, s}^{\Delta} \left(\sup_{k \geq 1} \|S_{1,k}(\varepsilon, \Delta, \mathcal{W}_M)\| \geq \delta \right) + \mathbb{P}_{\beta, s}^{\Delta} (\nu(\varepsilon) < \sigma(\mathcal{W}_M)) \right\} .$$

The convergence of the sequence $(s_k)_k$ follows from the proof of Theorem 5.5 of [6] which states the almost sure convergence due to the previous propositions. \square

REMARK 3. We can weaken the condition on p given in (A3'). Indeed, we can assume that (A3') holds for any $p > 0$ as soon as at least condition (4.6) is true also for power 2 of V . This is needed in the proof while giving an upper bound for all the T_n 's using the Jensen inequality instead of the Minkowski inequality as in [6]. In this case, the assumption (A4) would have to be satisfied for a power $\max(p, 2)$ instead of power 2.

4.2. Convergence Theorem for Dense Deformable Template Model. We now give the convergence result of our estimation process which is an application of the previous theorem. In this section, we assume that σ^2 is fixed which reduces θ to (α, Γ_g) . In fact, due to the implicit definition of $\hat{\theta}$ given in equation (3.1), we were not able to prove the smoothness of the inverse of the function $s \mapsto \hat{\theta}(s)$ which is straightforward for fixed σ^2 .

We can easily exhibit some of the functions involved in our procedure. Comparing equation (3.2) to equation (4.1), we have

$$H_s(\beta) = S(\beta) - s. \quad (4.16)$$

Equation (3.1) gives the existence of the function $s \rightarrow \hat{\theta}(s)$. We denote by l the observed log-likelihood : $l(\theta) \triangleq \log \int q(y, \beta, \theta) d\beta$, and let $w(s) \triangleq -l \circ \hat{\theta}(s)$ and $h(s) \triangleq \int H_s(\beta) q_{post}(\beta | y, \hat{\theta}(s)) d\beta$ for $s \in \mathcal{S}$.

THEOREM 4.4. The sequence of stochastic approximations $(s_k)_k$ related to the model defined in Section 2 and generated by Algorithms 1 and 2 satisfies the assumptions (A1') (ii), (iii), (iv), (A2) and (A3').

Proof. The details of the proof are given in appendix (Section 6). \square

COROLLARY 1 (Convergence of Dense Deformable Template building via Stochastic Approximation).

Assume

1. there exist $p \geq 1$ and $a \in]0, 1[$ such that the sequences $\Delta = (\Delta_k)_{k \geq 0}$ and $\varepsilon = (\varepsilon_k)_{k \geq 0}$ are non-increasing, positive and satisfy:

$$\sum_{k=0}^{\infty} \Delta_k = \infty, \quad \lim_{k \rightarrow \infty} \varepsilon_k = 0 \quad \text{and} \quad \sum_{k=1}^{\infty} \{\Delta_k^2 + \Delta_k \varepsilon_k^a + (\Delta_k \varepsilon_k^{-1})^p\} < \infty;$$
2. $\mathcal{L} \triangleq \{s \in \mathcal{S}, \langle \nabla w(s), h(s) \rangle = 0\}$ is included in a level set of w .

Let K be a compact subset of \mathbb{R}^N and K_0 a compact subset of $S(\mathbb{R}^N)$.

Let $(s_k)_{k \geq 0}$ and $(\theta_k)_{k \geq 0}$ be the two sequences defined in Algorithms 1 and 2. We denote by $\mathcal{L}' \triangleq \{\theta \in \hat{\theta}(\mathcal{S}), \frac{\partial l}{\partial \theta}(\theta) = 0\}$, then $\hat{\theta}(\mathcal{L}) = \mathcal{L}'$ and

$$\lim_{k \rightarrow \infty} d(\theta_k, \mathcal{L}') = 0 \quad \bar{\mathbb{P}}_{\beta_0, s_0, 0, 0, 0} - a.s.$$

for all $\beta_0 \in K$ and $s_0 \in K_0$, where $\bar{\mathbb{P}}_{\beta_0, s_0, 0, 0, 0}$ is the probability measure associated with the chain $(Z_k = (\beta_k, s_k, \kappa_k, \zeta_k, \nu_k))_{k \geq 0}$ starting at $(\beta_0, s_0, 0, 0, 0)$.

Proof. We first notice that, as mentioned in [8] (Lemma 2 equation 36), since $\hat{\theta}$, ϕ and ψ are smooth functions, it is easy to relate the convergence of the stochastic approximation sequence $(s_k)_k$ to the convergence of the estimated parameter sequence $(\theta_k)_k$.

Then the proof follows from the general stability result Theorem 4.1 stated in the subsection 4.1 and from the previous theorem 4.4. \square

REMARK 4. Note that condition (1) is easily checked for $\Delta_k = k^{-c}$ and $\varepsilon_k = k^{-c'}$ with $1/2 < c' < c < 1$. However, condition (2) has not been successfully proved yet and should be relaxed in future work.

5. Experiments. To illustrate our stochastic algorithm for the deformable template models, we consider handwritten digit images. For each digit class, we learn the template, the corresponding noise variance and the geometric covariance matrices. (Note that in this experiment the noise

variance is no longer fixed and is estimated as the other parameters). We use the US-Postal database which contains a training set of around 7000 images.

Each picture is a (16×16) gray level image with intensity in $[0, 2]$ where 0 corresponds to the black background. We will also use these sets in the special case of a noisy setting by adding independent centered Gaussian noise to each image.

To be able to compare the results with the previous deterministic algorithm proposed in [2], we use the same samples. In Figure 5.1 below, we show some of the training images.



FIG. 5.1. Some images from the training set used for the estimation of the model parameters (inverse video).



FIG. 5.2. Estimated prototypes of digit 1 (20 images per class) for different hyper-parameters. Left: smoother geometry but larger photometric covariance in the spline kernel. Right: more rigid geometry and smaller photometric covariance.



FIG. 5.3. Synthetic examples corresponding to the two previous estimated templates of digit 1 (inverse video). Left : with a fatty shape. Right : with a correct shape thickness.

A natural choice for the hyper-parameters on α and Γ_g is $\mu_p = 0$ and we induce the two covariance matrices Σ_p and Σ_g by the metric of the Hilbert spaces V_p and V_g (defined in Section 2.1) involving the correlation between the landmarks determined by the kernel. Define the square matrices

$$\begin{aligned} M_p(k, k') &= K_p(v_{p,j}, v_{p,j'}) \quad \forall 1 \leq k, k' \leq k_p \\ M_g(k, k') &= K_g(v_{g,j}, v_{g,j'}) \quad \forall 1 \leq k, k' \leq k_g, \end{aligned} \quad (5.1)$$

then $\Sigma_p = M_p^{-1}$ and $\Sigma_g = M_g^{-1}$. In our experiments, we have chosen Gaussian kernels for both K_p and K_g , where the standard deviations are fixed at $\sigma_p = 0.12$ and $\sigma_g = 0.3$. The deformation is computed in the $[-1, 1]^2$ square with $k_g = 6$ equi-distributed landmarks on this domain. The template has been estimated with $k_p = 15$ equi-distributed control points on $[-1.5, 1.5]^2$.

These two covariance matrices are important hyper-parameters; indeed, it has been shown in [2] that changing the geometric covariance has an effect on the sharpness of the template images. As for the photometric hyper-parameter, it affects both the template and the geometry in the sense that with a large variance, the kernel centered on one landmark spreads out to many of its neighbors. This leads to thicker shapes as shown in the left panel of Figure 5.2. As a consequence, the template is biased: it is not “centered” in the sense that the mean of the deformations required to fit the data is not close to zero. For example for digit “1”, the main deformations should be contractions or dilations of the template. With a large variance σ_p^2 , the template is thicker yielding



FIG. 5.4. *Estimated prototypes issued from left 10 images per class and right 20 images per class in the training set.*

larger contractions and smaller dilations. Since we have set a Gaussian law on the deformation variable β and $z_{-\beta} = -z_{\beta}$, the deformations $(Id + z_{\beta})$ and $(Id - z_{\beta})$ have the same probability to be drawn under the estimated model. As shown on synthetic examples given in Figure 5.3 left panel, there are many large dilated shapes. However, these examples were not in the training set and are not generated with other hyper-parameters (Figure 5.3 right panel). We have tried different relevant values and kept the best with regard to the visual results. We present in the following only the results with the adapted variances.

For the stochastic approximation step-size, we allow a heating period which corresponds to the absence of memory for the first iterations. This allows the Markov chain to reach a region of interest in the posterior probability density function before exploring this particular region.

In the experiments presented here, the heating time lasts k_h (up to 150) iterations and the whole algorithm stops after, at most, 200 iterations depending on the data set (noisy or not). This number of iterations corresponds to a point where the convergence seems to have been reached. This yields:

$$\Delta_k = \begin{cases} 1, & \forall 1 \leq k \leq k_h \\ \frac{1}{(k-k_h)^d}, & \forall k > k_h \text{ for } d = 0.6 \text{ or } 1. \end{cases}$$

To optimize the choice of the transition kernel Π_{θ} , we have run the algorithm with different kernels and compared the evolution of the simulated hidden variables as well as the results on the estimated parameters. Some kernels, as the ones mentioned above do not yield good coverage of the infinite support of the unobserved variable. From this point of view the hybrid Gibbs sampler we used has better properties and gives nice estimation results which are presented below.

5.1. Estimated Template. We show here the results of the statistical learning algorithm for this model. Figure 5.4 shows two runs of the algorithm for a non-noisy database with 10 and 20 images per class. Ten images per class are enough to obtain satisfactory template images with high contrast.

Although it was proved in [2] that the Kullback-Leibler divergence between $q(\cdot; \tilde{\theta})$ and the common density function for observations from a given class converges to its minimal value on the family $q(\cdot; \theta)$, we note that increasing the number of training images does not significantly improve the estimated photometric template. This apparently surprising fact can be explained as follows : since strong variations in appearance among the images may happen within a given class (think about topological changes for instance), the image distribution can not be perfectly represented as a distribution around a single template. This distribution is better represented as clustered around a major template and minor ones in a multimodal way. When the sample size is moderate, with a high probability, the sample contains basically images around the major mode and parametric model fits these data quite accurately. When the sample size increases, the minor modes start to play a significant role as “outliers” with respect to the major mode in the data, resulting in a slightly more blurry template trying to accommodate the different modes. One way to overcome this fact is to use some clustering methods as proposed in [2]. To visualize robustness with respect to the training set, we ran this algorithm with 20 images per class randomly chosen

in the whole database. The different runs are presented in Figure 5.5. The two left images show some templates which look like the ones obtained in the left image of Figure 5.4 with the 20 first examples of the database. When outliers appear among the 20 randomly chosen training images the template may become somewhat more blurry. This is observed for digits '2' and '4' (apparently the most variable digits) in the right image of Figure 5.5. For digits where the outliers are less far from the other images, the templates are stable.



FIG. 5.5. Templates estimated with randomly chosen samples from the whole USPostal database. Each image is one run of the algorithm with same initial conditions but different training sets of 20 images per digit each. The variability of the results is related to the huge variability inside the USPS database.




FIG. 5.6. Evolution of the templates with the algorithm iterations. Top line. Left: Mean gray level images of the 20 training samples. Middle: template at the 50th iteration. Right: template at the 100th iteration. Bottom line: template at the 150th iteration. The improvement is visible, very fast for some very simple shapes as digit 1 and longer for very variable ones as digit 2. The higher geometrical variability increases the fitting time of the algorithm.

The evolution of the template with the iterations can be viewed in Figure 5.6. The initialization of the template is the mean of the gray level images. As the iterations proceed, the templates become sharper. In particular, the estimated templates for digits with small geometrical variability converge very fast. For digits like '2' or '4', where the geometrical variability is higher, the convergence of the coupled parameters (photometry and geometry) is slowed down.

5.2. Photometric noise variance. The evolution of the noise variance all along the SAEM-MCMC iterations is the same as the one observed with the “mode approximation EM” described in [2]. During the first iterations, the noise variance balances the inaccuracy of the estimated template which is simply the gray-level mean of the training set. As the iterations proceed, the template estimates become sharper as does the estimate of the covariance matrix for the geometry.




FIG. 5.7. Evolution of the estimated noise variance using 20 images per class along the SAEM-MCMC algorithm. This confirms the visual effects seen on the templates : rapid convergence for some really constrained digits and slower convergence for very variable ones.

This yields very small residual noise. Note that here the final noise variance, which is less than 0.1, for the SAEM-MCMC algorithm for all digits is less than the noise variance, which is between 0.2 and 0.3, for the mode approximation EM experimented in [2] in the one component run. This can be explained by the stochastic nature of the algorithm which enables it to escape from local minima provoking early terminations in the deterministic version.

5.3. Estimated geometric distribution. As mentioned previously we have to fix the value of the hyper-parameter a_g of the prior on Γ_g . This quantity plays a significant role in the results. Indeed, to satisfy the theoretical conditions we have to choose a_g larger than $4k_g + 1$ say $4 \times 36 + 1$ in our examples. From the geometry update equation, a barycenter between the ‘sample’ covariance and the prior, with the number n of images and a_g as coefficients, we find that the prior dominates when the training set is small. The covariance matrix stays close to the prior. Thus we need to decrease a_g and find the best trade-off between the degenerate inverse Wishart and the weight of the prior in the covariance estimation. We fix this value with a visual criterion: both the templates and the generated sample with the learnt geometry have to be satisfactory. This yields $a_g = 0.5$ or 0.1.

As we have observed from Figure 5.9, parameter estimation is robust regardless of whether the prior is degenerate or not. In addition, considering the update formulas, even if this law does not have a total weight equal to 1 it does not affect parameter estimation.

In Figure 5.9, we show a sample of some synthetic digits modeled by deformation templates drawn with the estimated parameters. Note that the resulting digits in Figure 5.9 look like some elements of the training set and seem to explain these data correctly, whereas the prior produces some non-relevant local deformations (cf. Figure (5.8)). In particular, for some especially geometrically constrained digits such as 0 or 1, the geometry variability reflects their constraints. For digits like the 2s, the training set is heterogeneous and shows a large geometrical variability. When comparing to the deformations obtained by the mode approximation to EM in [2], it seems that here we obtain a more variable geometry. This might be because with a stochastic algorithm, we explore the posterior density and do not only concentrate at its mode. This allows some more exotic deformations corresponding to realizations of the missing variable β which may belong to




FIG. 5.8. *Effect of the prior distribution on the deformation : 20 synthetic examples per class generated with the estimated template but the prior covariance matrix (inverse video).*



FIG. 5.9. *Effect of the estimated geometric distribution : 40 synthetic examples per class generated with the estimated parameters: 20 with the direct deformations and 20 with the symmetric deformations (inverse video).*

the tail of the law. Another reason may be that for such digits, the mode approximation gets stuck in a local minimum of the matching energy. Jumping out of this configuration would require a large deformation (not allowed by the gradient descent since it would increase the energy again). However, such a deformation can be proposed leading to acceptance by the stochastic algorithm. Subsequently the deformed template may better fit the observations, leading to acceptance of these large deformations. This also leads to a lower value of the residual noise and may also explain the low noise variance estimated by the stochastic EM algorithm.

5.4. Noise effect. As shown in [2], in the presence of noise, the mode approximation algorithm does not converge towards the MAP estimator. In our setting, the consistency of the “SAEM like” algorithm has been proved independently of the training set, and thus noisy images can also be treated exactly the same way. These are the results we present here. Figure 5.10 shows two training examples per class for noise variance values $\sigma^2 = 1$ and $\sigma^2 = 2$. In Figures 5.11 and




FIG. 5.10. Two images examples per class of the noisy training set (variance: top: $\sigma^2 = 1$, bottom: $\sigma^2 = 2$).



FIG. 5.11. Estimated prototypes in a noisy setting $\sigma^2 = 1$: Left: with the mode approximation algorithm. Right: with the SAEM-MCMC coupling procedure.

5.12, we show the estimated templates for the noisy training set containing 20 images for both methods. Even if the mode approximation algorithm does not diverge, it cannot fit the template for digits with a high variability. In contrast, the stochastic EM gives acceptable contrasted templates which look like those obtained in Figure 5.4. This becomes more significant as we increase the variance of the additive noise we introduce in the training set.

Concerning the choice of the hyper-parameters, it is not necessary to change all of them. For the photometric variance of the spline kernel, a small one could create some non-smooth templates and a large kernel would smooth the noise effect. However, we can keep the geometric hyper-parameters unchanged. We are presenting here only experiments which seemed to provide a reasonable tradeoff between these effects.

The geometry is also well estimated despite the high level of noise in the training set. Figure 5.13 shows some synthetic examples, in which parameters are learnt from the training set with an additive noise variance of one. The two lines correspond to deformations and their symmetric deformation. This sample looks like the synthetic samples learnt on non-noisy images even if some examples are not relevant. However, the global behavior has been learnt.

The algorithm manages to catch the photometry (a contrasted and smoothed template), the geometry of the shapes and to “separate” the additive noise.

The number of iterations needed to reach the convergence point in the noisy setting is about twice that of the non-noisy case. The template takes the longest time to converge and the estimate of σ^2 converges in a few iterations. In particular, the templates obtained in the left panel of Figure 5.4 with only 10 images per training digit set are obtained with a heating period of 25 iterations and 5 more steps with memory. The templates of Figure 5.11, right picture, require 100 to 125 heating iterations in the 150 global iterations. This is understandable since the algorithm has to cope with variations due to the noise and thus needs a longer time to fit the model.




FIG. 5.12. *Estimated prototypes in a noisy setting $\sigma^2 = 2$: Left: with the mode approximation algorithm. Right: with the SAEM-MCMC coupling procedure.*



FIG. 5.13. *Effect of the noise on the geometric parameter estimation : 40 synthetic examples per class generated with the parameters estimated from the noisy training set (additive noise variance of 1, inverse video).*

6. Proof of Theorem 4.4 . Here we demonstrate Theorem 4.4, i.e. the stochastic approximation sequence satisfies assumptions (A1') (ii), (iii), (iv), (A2) and (A3').

We recall that in this section, the parameter σ^2 is fixed so that $\theta = (\alpha, \Gamma)$. The sufficient statistic vector S , the set \mathcal{S} as well as the explicit expression of $\hat{\theta}(s)$ have been given in Subsection 4.2. As noted, $\hat{\theta}$ is a smooth function of \mathcal{S} .

We will prove that these conditions hold for any $p \geq 1$ and $a \in]0, 1[$.

6.1. Proof of assumption (A1').

We recall the functions H , h and w as in [8] defined as follows:

$$\begin{aligned} H_s(\beta) &= S(\beta) - s, \\ h(s) &= \int_{\mathbb{R}^N} H_s(\beta) q_{post}(\beta | \mathbf{y}, \hat{\theta}(s)) d\beta, \\ w(s) &= -l(\hat{\theta}(s)). \end{aligned}$$

As shown in [8], with these functions, we satisfy (A1'(iii)) and (A1'(iv)).

Moreover, since the interpolation kernel K_p is bounded, there exist $A > 0$ and $B \in \text{Sym}_{k_p}^+$ such that for any $\beta \in \mathbb{R}^N$, we have

$$\|S_1(\beta)\| \leq A, \quad 0 \leq S_2(\beta) \leq B \quad \text{and} \quad 0 \leq S_3(\beta),$$

where, for any symmetric matrices B and B' , we say that $B \leq B'$ if $B' - B$ is a non-negative symmetric matrix.

We define the set \mathcal{S}_a by

$$\mathcal{S}_a \triangleq \{ S \in \mathcal{S} \mid \|S_1\| \leq A, \quad 0 \leq S_2 \leq B \quad \text{and} \quad 0 \leq S_3 \}.$$

Since the constraints are obviously convex and closed, we get that \mathcal{S}_a is a closed convex subset of \mathbb{R}^{n_s} such that

$$\mathcal{S}_a \subset \mathcal{S} \subset \mathbb{R}^{n_s}$$

and satisfying

$$s + \rho H_s(\beta) \in \mathcal{S}_a \quad \text{for any } \rho \in [0, 1] \text{ any } s \in \mathcal{S}_a \text{ and any } \beta \in \mathbb{R}^N.$$

We now focus on the first two points. As l and $\hat{\theta}$ are continuous functions, we only need to prove that $\mathcal{W}_M \cap \mathcal{S}_a$ is a bounded set for a constant $M \in \mathbb{R}_+^*$ with:

$$\mathcal{W}_M = \{s \in \mathcal{S}, \quad w(s) \leq M\}.$$

On \mathcal{S}_a , s_1 and s_2 are bounded; writing $\hat{\theta}(s) = (\alpha(s), \Gamma(s))$, we deduce from (3.1) and from the boundedness of K_p that $\alpha(s)$ is bounded on \mathcal{S}_a and $|y_i - K_p^{\beta_i} \alpha(s)|$ is uniformly bounded on $\beta_i \in \mathbb{R}^{2k_g}$ and $s \in \mathcal{S}_a$. Hence (recall that σ^2 is fixed here), there exists an $\eta > 0$ such that $q_c(\mathbf{y}|\beta, \hat{\theta}(s)) \geq \eta$ for any $s \in \mathcal{S}_a$ and $\beta \in \mathbb{R}^N$. Thus,

$$w(s) \geq -\log \left(\int q_m(\beta, \hat{\theta}(s)) d\beta \right) + C \geq -\log(q_{para}(\hat{\theta}(s))) + C \geq -\log(q_{para}|_{\Gamma}(\Gamma(s))) + C,$$

where C is a constant independent of $s \in \mathcal{S}_a$. Since

$$-\log(q_{para}|_{\Gamma}(\Gamma_g)) = \frac{a_g}{2} (\langle \Gamma_g^{-1}, \Sigma_g \rangle_F + \log |\Gamma_g|) \geq \frac{a_g}{2} \log |\Gamma_g|$$

and

$$\lim_{\|s\| \rightarrow +\infty, s \in \mathcal{S}_a} \log(|\Gamma_g(s)|) = \lim_{\|s\| \rightarrow +\infty, s \in \mathcal{S}_a} \log(|(s_3 + a_g \Sigma_g)/(n + a_g)|) = +\infty,$$

we deduce that

$$\lim_{\|s\| \rightarrow +\infty, s \in \mathcal{S}_a} w(s) = +\infty.$$

Since w is continuous and \mathcal{S}_a is closed, this proves **(A1'ii)**.

6.2. Proof of assumption (A2).

We prove a classical sufficient condition **(DRI1)**, used in [6] which will imply **(A2)** under the condition that H_s is dominated by V for any $s \in \mathcal{K}$.

(DRI1) For any $s \in \mathcal{S}$, $\Pi_{\hat{\theta}(s)}$ is ϕ -irreducible and aperiodic. In addition there exist a function $V : \mathbb{R}^N \rightarrow [1, \infty[$ and $p \geq 2$ such that for any compact subset $\mathcal{K} \subset \mathcal{S}$, there exist an integer m and constants $0 < \lambda < 1$, $B > 0$, $\kappa > 0$, $\delta > 0$, a subset \mathcal{C} of \mathbb{R}^N and a probability measure ν such that

$$\sup_{s \in \mathcal{K}} \Pi_{\hat{\theta}(s)}^m V^p(\beta) \leq \lambda V^p(\beta) + B \mathbf{1}_{\mathcal{C}}(\beta), \quad (6.1)$$

$$\sup_{s \in \mathcal{K}} \Pi_{\hat{\theta}(s)} V^p(\beta) \leq \kappa V^p(\beta) \quad \forall \beta \in \mathbb{R}^N, \quad (6.2)$$

$$\inf_{s \in \mathcal{K}} \Pi_{\hat{\theta}(s)}^m(\beta, A) \geq \delta \nu(A) \quad \forall \beta \in \mathcal{C}, \forall A \in \mathcal{B}(\mathbb{R}^N). \quad (6.3)$$

REMARK 5. Note that condition (6.3) is equivalent to the existence of a small set \mathcal{C} (defined below) which only depends on \mathcal{K} .

NOTATION 1. Let $(e_j)_{1 \leq j \leq N}$ be the canonical basis of \mathbb{R}^N . For any $1 \leq j \leq N$, let $E_{\theta,j} \triangleq \{ \beta \in \mathbb{R}^N \mid \langle \beta, e_j \rangle_\theta = 0 \}$ be the orthogonal space of $\text{Span}\{e_j\}$ and $p_{\theta,j}$ be the orthogonal projection on $E_{\theta,j}$ i.e.

$$p_{\theta,j}(\beta) \triangleq \beta - \frac{\langle \beta, e_j \rangle_\theta}{\|e_j\|_\theta^2} e_j,$$

where $\langle \beta, \beta' \rangle_\theta = \sum_{i=1}^n \beta_i^t \Gamma_g^{-1} \beta'_i$ for $\theta = (\alpha, \Gamma_g)$ (i.e. the natural dot product associated with the covariance matrix Γ_g) and $\|\cdot\|_\theta$ the corresponding norm.

We denote for any $1 \leq j \leq N$ and $\theta \in \Theta$ by $\Pi_{\theta,j}$ the Markov kernel on \mathbb{R}^N (3.3) associated with the Metropolis-Hastings step of the j -th Gibbs sampler step on β . We have $\Pi_\theta = \Pi_{\theta,N} \circ \cdots \circ \Pi_{\theta,1}$.

We first recall the definition of a small set:

DEFINITION 1. (cf. [15]) A set $\mathcal{E} \in \mathcal{B}(\mathcal{X})$ is called a **small set** for the kernel Π if there exist an $m > 0$, and a non trivial measure ν_m on $\mathcal{B}(\mathcal{X})$, such that for all $\beta \in \mathcal{E}$, $B \in \mathcal{B}(\mathcal{X})$,

$$\Pi^m(\beta, B) \geq \nu_m(B). \quad (6.4)$$

When (6.4) holds, we say that \mathcal{E} is ν_m -small.

We now prove the following lemma which give the existence of the small set \mathcal{C} in (DRI1):

LEMMA 6.1. Let \mathcal{E} be a compact subset of \mathbb{R}^N and \mathcal{K} a compact subset of \mathcal{S} . Then \mathcal{E} is a small set of \mathbb{R}^N for $\Pi_{\hat{\theta}(s)}$ for any $s \in \mathcal{K}$.

Proof. First note that there exists an $a_c > 0$ such that for any $\theta \in \Theta$, any $\beta \in \mathbb{R}^N$ and any $b \in \mathbb{R}$, the acceptance rate $r_j(\beta^j, b; \beta^{-j}, \theta)$ is uniformly bounded below by a_c so that for any $1 \leq j \leq N$ and any non-negative function f ,

$$\Pi_{\theta,j} f(\beta) \geq a_c \int_{\mathbb{R}} f(\beta^{-j} + b e_j) q_j(b | \beta^{-j}, \theta) db = a_c \int_{\mathbb{R}} f(p_{\theta,j}(\beta) + z e_j / \|e_j\|_\theta) g_{0,1}(z) dz,$$

where $g_{0,1}$ is the density of the standard Gaussian distribution $\mathcal{N}(0, 1)$.

By induction, we have

$$\Pi_\theta f(\beta) \geq a_c^N \int_{\mathbb{R}^N} f \left(p_{\theta,N,1}(\beta) + \sum_{j=1}^N z_j p_{\theta,N,j+1}(e_j) / \|e_j\|_\theta \right) \prod_{j=1}^N g_{0,1}(z_j) dz_j, \quad (6.5)$$

where $p_{\theta,q,r} = p_{\theta,r} \circ p_{\theta,r-1} \circ \cdots \circ p_{\theta,q}$ for any integers $q \leq r$ and $p_{\theta,N,N+1} = \text{Id}$.

Let $A_\theta \in \mathcal{L}(\mathbb{R}^N)$ be the linear mapping on \mathbb{R}^N defined by

$$A_\theta z = \sum_{j=1}^N z_j p_{\theta,N,j+1}(e_j) / \|e_j\|_\theta.$$

One easily checks that for any $1 \leq k \leq N$, $\text{Span}\{ p_{\theta,N,j+1}(e_j), k \leq j \leq N \} = \text{Span}\{e_j \mid k \leq j \leq N\}$ so that A_θ is an invertible mapping. By a change of variable, we get

$$\int_{\mathbb{R}^N} f(p_{\theta,N,1}(\beta) + A_\theta z_1^N) \prod_{j=1}^N g_{0,1}(z_j) dz_j = \int_{\mathbb{R}^N} f(u) g_{p_{\theta,N,1}(\beta), A_\theta A_\theta^t}(u) du,$$

where $g_{\mu,\Sigma}$ stands for the density of the normal law $\mathcal{N}(\mu, \Sigma)$. Since $\theta \rightarrow A_\theta$ is smooth on the set of invertible mappings in θ , we deduce that there exist two constants $c_K > 0$ and $C_K > 0$ such that $c_K \text{Id} \leq A_\theta A_\theta^t \leq \text{Id} / c_K$ and $g_{p_{\theta,N,1}(\beta), A_\theta A_\theta^t}(u) \geq C_K g_{p_{\theta,N,1}(\beta), \text{Id} / c_K}(u)$ uniformly for $\theta = \hat{\theta}(s)$ with $s \in \mathcal{K}$. Assuming that $\beta \in \mathcal{E}$, since $\theta \rightarrow p_{\theta,N,1}$ is smooth and \mathcal{E} is compact, we have

$\sup_{\beta \in \mathcal{E}, \theta = \hat{\theta}(s), s \in \mathcal{K}} \|p_{\theta, N, 1}(\beta)\| < \infty$. Therefore, there exist $C'_\mathcal{K} > 0$ and $c'_\mathcal{K} > 0$ such that for any $(u, \beta) \in \mathbb{R}^N \times \mathcal{E}$ and any $\theta = \hat{\theta}(s)$, $s \in \mathcal{K}$

$$g_{p_{\theta, N, 1}(\beta), A_\theta A_\theta^t}(u) \geq C'_\mathcal{K} g_{0, \text{Id}/c'_\mathcal{K}}(u). \quad (6.6)$$

Using (6.5) and (6.6), we deduce that for any A , for any $s \in \mathcal{K}$ and $\theta = \hat{\theta}(s)$,

$$\Pi_\theta(\beta, A) \geq C'_\mathcal{K} a_c^N \nu_\mathcal{K}(A),$$

with $\nu_\mathcal{K}$ equals to the density of the normal law $\mathcal{N}(0, \text{Id}/c'_\mathcal{K})$.

This yields the existence of the small set as well as equation (6.3). \square

This property also implies the ϕ -irreducibility of the Markov chain $(\beta_k)_k$ and its aperiodicity (cf. [15] p121).

We set $V : \mathbb{R}^N \rightarrow [1, +\infty[$ as the following function

$$V(\beta) = 1 + \|\beta\|^2. \quad (6.7)$$

We, in fact, have the following property : $\exists C_\mathcal{K} > 0$ such that : $\forall \beta \in \mathbb{R}^N$,

$$\sup_{s \in \mathcal{K}} \|H_s(\beta)\| \leq C_\mathcal{K} V(\beta).$$

This condition is required for the implication of **(A2)** by **(DRI1)**.

We now prove condition (6.2).

Let \mathcal{K} be a compact subset of \mathcal{S} and $p \geq 1$. For any $1 \leq j \leq N$, any $s \in \mathcal{K}$ and $\theta = \hat{\theta}(s)$, we have

$$\Pi_{\theta, j} V^p(\beta) \leq V^p(\beta) + \int_{\mathbb{R}} V^p(p_{\theta, j}(\beta) + ze_j/\|e_j\|_\theta) g_{0, 1}(z) dz.$$

Since $V(\beta + h) \leq 2(V(\beta) + V(h))$ for any $\beta, h \in \mathbb{R}^N$ and since there exist two constants $c_\mathcal{K} > 0$ and $C_\mathcal{K} > 0$ such that for any $\beta \in \mathbb{R}^N$, $\theta \in \hat{\theta}(\mathcal{K})$, $\|p_{\theta, j}(\beta)\| \leq C_\mathcal{K} \|\beta\|$ and $\|e_j\|_\theta \geq 1/c_\mathcal{K}$, we have

$$\int_{\mathbb{R}} V^p(p_{\theta, j}(\beta) + ze_j/\|e_j\|_\theta) g_{0, 1}(z) dz \leq 2^p C_\mathcal{K}^p V^p(\beta) \int_{\mathbb{R}} (1 + V(c_\mathcal{K} z e_j))^p g_{0, 1}(z) dz.$$

We deduce that there exists an $C'_\mathcal{K} > 0$ such that for any $\beta \in \mathbb{R}^N$

$$\sup_{\theta = \hat{\theta}(s), s \in \mathcal{K}} \Pi_{\theta, j} V^p(\beta) \leq C'_\mathcal{K} V^p(\beta).$$

Then, by composition $\Pi_\theta V^p(\beta) \leq C'^N_\mathcal{K} V^p(\beta)$ and (6.2) holds for any $p \geq 1$.

Now consider the Drift condition (6.1).

To prove this inequality, we prove the same inequality for a subsidiary function V_θ which depends on the parameters θ and then we deduce the result for V . So let us define for any $\theta = (\alpha, \Gamma_g)$ the function $V_\theta(\beta) \triangleq 1 + \|\beta\|_\theta^2$.

LEMMA 6.2. *Let K be a compact subset of Θ . For any $p \geq 1$, there exist an $0 \leq \rho_K < 1$ and an $C_K > 0$ such that for any $\theta \in K$, any $\beta \in \mathbb{R}^N$ we have*

$$\Pi_\theta V_\theta^p(\beta) \leq \rho_K V_\theta^p(\beta) + C_K.$$

Proof. The proposal distribution for $\Pi_{\theta,j}$ is given by $q(\beta \mid \beta^{-j}, y, \theta) \stackrel{\text{law}}{=} p_{\theta,j}(\beta) + z \frac{e_j}{\|e_j\|_\theta}$ where $z \sim \mathcal{N}(0, 1)$. Then, there exists C_K such that for any $\beta \in \mathbb{R}^N$ and any measurable set $A \in \mathcal{B}(\mathbb{R}^N)$

$$\Pi_{\theta,j}(\beta, A) = (1 - a_{\theta,\beta}) \mathbb{1}_A(\beta) + a_{\theta,\beta} \int_{\mathbb{R}} \mathbb{1}_A \left(p_{\theta,j}(\beta) + z \frac{e_j}{\|e_j\|_\theta} \right) g_{0,1}(z) dz,$$

where $a_{\theta,\beta} \geq a_c$ (a_c is a lower bound for the acceptance rate),

Since $\langle p_{\theta,j}(\beta), e_j \rangle_\theta = 0$, we get $V_\theta \left(p_{\theta,j}(\beta) + z \frac{e_j}{\|e_j\|_\theta} \right) = V_\theta(p_{\theta,j}(\beta)) + z^2$ and

$$\begin{aligned} \Pi_{\theta,j} V_\theta^p(\beta) &= (1 - a_{\theta,\beta}) V_\theta^p(\beta) + a_{\theta,\beta} \int_{\mathbb{R}} (V_\theta(p_{\theta,j}(\beta)) + z^2)^p g_{0,1}(z) dz \\ &\leq (1 - a_{\theta,\beta}) V_\theta^p(\beta) + a_{\theta,\beta} \left(V_\theta^p(p_{\theta,j}(\beta)) + C_K V_\theta^{p-1}(p_{\theta,j}(\beta)) \int_{\mathbb{R}} (1 + z^2)^p g_{0,1}(z) dz \right) \\ &\leq (1 - a_{\theta,\beta}) V_\theta^p(\beta) + a_{\theta,\beta} V_\theta^p(p_{\theta,j}(\beta)) + C'_K V_\theta^{p-1}(p_{\theta,j}(\beta)). \end{aligned}$$

We have used in the last inequality the fact that a Gaussian variable has bounded moments of any order. Since $a_{\theta,\beta} \geq a_c$ and $\|p_{\theta,j}(\beta)\|_\theta \leq \|\beta\|_\theta$ ($p_{\theta,j}$ is an orthonormal projection for the dot product $\langle \cdot, \cdot \rangle_\theta$), we get that $\forall \eta > 0$, $\exists C_{K,\eta}$ such that $\forall \beta \in \mathbb{R}^N$ and $\forall \theta \in K$

$$\Pi_{\theta,j} V_\theta^p(\beta) \leq (1 - a_c) V_\theta^p(\beta) + (a_c + \eta) V_\theta^p(p_{\theta,j}(\beta)) + C_{K,\eta}.$$

By induction, we show that

$$\Pi_\theta V_\theta^p(\beta) \leq \sum_{u \in \{0,1\}^N} \prod_{j=1}^N (1 - a_c)^{1-u_j} (a_c + \eta)^{u_j} V_\theta^p(p_{\theta,u}(\beta)) + \frac{C_{K,\eta}}{\eta} ((1 + \eta)^{N+1} - 1),$$

where $p_{\theta,u} = ((1 - u_N) \text{Id} + u_N p_{\theta,N}) \circ \dots \circ ((1 - u_1) \text{Id} + u_1 p_{\theta,1})$. Let $p_\theta = p_{\theta,N} \circ \dots \circ p_{\theta,1}$ and note that $p_{\theta,j}$ is contracting so that

$$\Pi_\theta V_\theta^p(\beta) \leq b_{c,\eta} V_\theta^p(\beta) + (a_c + \eta)^N V_\theta^p(p_\theta(\beta)) + \frac{C_{K,\eta}}{\eta} ((1 + \eta)^{N+1}),$$

for $b_{c,\eta} = \left(\sum_{u \in \{0,1\}^N, u \neq \mathbf{1}} \prod_{j=1}^N (1 - a_c)^{1-u_j} (a_c + \eta)^{u_j} \right)$.

To end the proof, we need to check that p_θ is strictly contracting uniformly on K . Indeed, $\|p_\theta(\beta)\|_\theta = \|\beta\|_\theta$ implies that $p_{\theta,j}(\beta) = \beta$ for any $1 \leq j \leq N$. This yields $\langle \beta, e_j \rangle_\theta = 0$ and thus $\beta = 0$ since $(e_j)_{1 \leq j \leq N}$ is a basis. Using the continuity of the norm of p_θ in θ and the compactness of K , we deduce that there exists $0 < \rho_K < 1$ such that $\|p_\theta(\beta)\|_\theta \leq \rho_K \|\beta\|_\theta$ for any β and $\theta \in K$. Changing ρ_K for $1 > \rho'_K > \rho_K$ we get $(1 + \rho_K^2 \|\beta\|_\theta^2)^p \leq \rho_K'^{2p} (1 + \|\beta\|_\theta^2)^p + C_K''$ for some uniform constant C_K'' . Therefore,

$$\Pi_\theta V_\theta^p(\beta) \leq b_{c,\eta} V_\theta^p(\beta) + \rho_K'^{2p} (a_c + \eta)^N V_\theta^p(\beta) + C_{K,\eta}''.$$

Since we have $\inf_{\eta > 0} b_{c,\eta} + \rho_K'^{2p} (a_c + \eta)^N < 1$ the result is immediate. \square

Next, we prove the expected inequality for the function V .

LEMMA 6.3. *For any compact set $K \subset \Theta$, any $p \geq 1$, there exist $0 < \rho_K < 1$, $C_K > 0$ and m_0 such that $\forall m \geq m_0$, $\forall \theta \in K$, $\forall \beta \in \mathbb{R}^N$*

$$\Pi_\theta^m V^p(\beta) \leq \rho_K V^p(\beta) + C_K.$$

Proof. Indeed, there exist $0 \leq c_1 \leq c_2$ such that $c_1 V(\beta) \leq V_\theta(\beta) \leq c_2 V(\beta)$ for any $(\beta, \theta) \in \mathbb{R}^N \times K$. Then, using the previous lemma, we have $\Pi_\theta^m V^p(\beta) \leq c_1^{-p} \Pi_\theta^m V_\theta^p(\beta) \leq c_1^{-p} (\rho_K^m V_\theta^p(\beta) + C_K / (1 - \rho_K)) \leq (c_2 / c_1)^p (\rho_K^m V^p(\beta) + C_K / (1 - \rho_K))$. Choosing m large enough for $(c_2 / c_1)^p \rho_K^m < 1$ gives the result. \square

This finishes the proof of (6.1) and at the same time of (A2).

6.3. Proof of assumption (A3').

The geometric ergodicity of the Markov chain, implied by the Drift condition (6.1), ensures the existence of a solution of the Poisson equation (cf. [15]):

$$g_{\hat{\theta}(s)}(\beta) = \sum_{k \geq 0} (\Pi_{\hat{\theta}(s)}^k H_s(\beta) - h(s)).$$

We first prove condition (A3'(i)).

Since $H_s(\beta) = S(\beta) - s$ with $S(\beta)$ at most quadratic in β , the choice of V directly ensures (4.3).

Due to the result presented in [10], there exist upper bounds for the convergence rates and the constants involved in the quantification of the geometrical ergodicity of all the chains indexed by $s \in \mathcal{K}$ which only depend on m, λ, B, δ . Therefore, these constants only depend on the fixed compact set \mathcal{K} . This yields the uniform ergodicity of the family of Markov chains on \mathcal{K} . So there exist constants $0 < \gamma_{\mathcal{K}} < 1$ and $C_{\mathcal{K}} > 0$ such that

$$\|g_{\hat{\theta}(s)}\|_V = \left\| \sum_{k \geq 0} (\Pi_{\hat{\theta}(s)}^k H_s(\beta) - h(s)) \right\|_V \leq \sum_{k \geq 0} C_{\mathcal{K}} \gamma_{\mathcal{K}}^k \|H_s\|_V < \infty.$$

Thus $\forall s \in \mathcal{K}$, $g_{\hat{\theta}(s)}$ belongs to $\mathcal{L}_V = \{g : \mathbb{R}^N \rightarrow \mathbb{R}, \|g\|_V < \infty\}$.

Repeating the same calculation as above, it is immediate that $\Pi_{\hat{\theta}(s)} g_{\hat{\theta}(s)}$ belongs to \mathcal{L}_V too. This ends the proof of (A3'(i)).

We now move to the Hölder condition (A3'(ii)). We will use the following lemmas which state Lipschitz conditions on the transition kernel and its iterates:

LEMMA 6.4. *Let \mathcal{K} be a compact subset of \mathcal{S} . There exists a constant $C_{\mathcal{K}}$ such that for any $p \geq 1$ and any function $f \in \mathcal{L}_{V^p}$, $\forall (s, s') \in \mathcal{K}^2$ we have :*

$$\|\Pi_{\hat{\theta}(s)} f - \Pi_{\hat{\theta}(s')} f\|_{V^{p+1/2}} \leq C_{\mathcal{K}} \|f\|_{V^p} \|s - s'\|.$$

Proof. For any $1 \leq j \leq N$ and $f \in \mathcal{L}_{V^p}$, we have

$$\Pi_{\theta,j} f(\beta) = (1 - r_j(\beta, \theta)) f(\beta) + \int_{\mathbb{R}} f(\beta_{b \rightarrow j}) r_j(\beta^j, b; \beta^{-j}, \theta) q_j(b | \beta^{-j}, \theta) db,$$

where $r_j(\beta, \theta) = \int_{\mathbb{R}} r_j(\beta^j, b; \beta^{-j}, \theta) q_j(b | \beta^{-j}, \theta) db$ is the average acceptance rate.

Let s and s' be two points in \mathcal{K} and $s(\epsilon) = (1 - \epsilon)s + \epsilon s'$ for $\epsilon \in [0, 1]$ be a linear interpolation between s and s' (since \mathcal{S} is convex, we can assume that \mathcal{K} is a convex set so that $s(\epsilon) \in \mathcal{K}$ for any $\epsilon \in [0, 1]$). We denote also by $\theta(\epsilon) \triangleq \hat{\theta}(s(\epsilon))$ the associated path in Θ which is a continuously differentiable function. To study the difference $\|(\Pi_{\theta(1),j} - \Pi_{\theta(0),j}) f(\beta)\|$, introduce $\Pi_{\theta,j}^1 f(\beta) \triangleq (1 - r_j(\beta, \theta)) f(\beta)$ and $\Pi_{\theta,j}^2 f(\beta) \triangleq \int_{\mathbb{R}} f(\beta_{b \rightarrow j}) r_j(\beta^j, b; \beta^{-j}, \theta) q_j(b | \beta^{-j}, \theta) db$. We start with the difference $\|(\Pi_{\theta(1),j}^2 - \Pi_{\theta(0),j}^2) f(\beta)\|$. First note that under the conditional law $q_j(b | \beta^{-j}, \theta)$, $b \sim \mathcal{N}(b_{\theta,j}(\beta), 1/\|e_j\|_{\theta}^2)$ where

$$b_{\theta,j}(\beta) \triangleq e_j^t p_{\theta,j}(\beta) = e_j^t \beta - \langle \beta, e_j \rangle_{\theta} / \|e_j\|_{\theta}^2$$

is the j -th coordinate of $p_{\theta,j}(\beta)$. We have

$$\Pi_{\theta,j}^2 f(\beta) = \int_{\mathbb{R}} f(\beta_{0 \rightarrow j} + b e_j) r_j(\beta^j, b; \beta^{-j}, \theta) \exp\left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_{\theta}^2}{2}\right) \frac{\|e_j\|_{\theta}}{\sqrt{2\pi}} db.$$

Since $r_j(\beta^j, b; \beta^{-j}, \theta) = \tilde{r}_j(\beta^j, b; \beta^{-j}, \theta) \wedge 1$ where $\tilde{r}_j(\beta^j, b; \beta^{-j}, \theta) \triangleq \frac{q_{obs}(\mathbf{y} | \beta_{b \rightarrow j}, \theta)}{q_{obs}(\mathbf{y} | \beta, \theta)}$ is a smooth function in θ , we have

$$\|(\Pi_{\theta(1),j}^2 - \Pi_{\theta(0),j}^2) f(\beta)\| \leq \int_0^1 \int_{\mathbb{R}} \|f(\beta_{0 \rightarrow j} + b e_j)\| \left| \frac{d}{d\epsilon} \left(r_j(\beta^j, b; \beta^{-j}, \theta) \exp\left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_{\theta}^2}{2}\right) \frac{\|e_j\|_{\theta}}{\sqrt{2\pi}} \right) \right| db d\epsilon.$$

However, one easily checks that there exists a constant C_K such that for any $s, s' \in \mathcal{K}$, ϵ, j and β (with $\theta = \theta(\epsilon)$):

$$\begin{aligned} & \left| \frac{d}{d\epsilon} \exp \left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_\theta^2}{2} \right) \frac{\|e_j\|_\theta}{\sqrt{2\pi}} \right| \\ & \leq C_K (1 + |b - b_{\theta,j}(\beta)|)^2 \exp \left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_\theta^2}{2} \right) \frac{\|e_j\|_\theta}{\sqrt{2\pi}} \left(\left| \frac{d}{d\epsilon} b_{\theta,j}(\beta) \right| + \left| \frac{d}{d\epsilon} \|e_j\|_\theta \right| \right). \end{aligned} \quad (6.8)$$

Since $\frac{d}{d\epsilon} \|e_j\|_\theta = \frac{1}{2\|e_j\|_\theta} e_j^t \frac{d}{d\epsilon} \Gamma_\theta^{-1} e_j$, $\frac{d}{d\epsilon} \Gamma_\theta^{-1} = -\Gamma_\theta^{-1} \frac{d}{d\epsilon} \Gamma_\theta \Gamma_\theta^{-1}$ and $\frac{d}{d\epsilon} \Gamma_\theta = \frac{s'_3 - s_3}{n + a_g}$ (see (3.1)), we deduce that there exists another constant C_K such that

$$\left| \frac{d}{d\epsilon} \|e_j\|_\theta \right| \leq C_K \|s' - s\|. \quad (6.9)$$

Similarly, updating the constant C_K , we have¹

$$\left| \frac{d}{d\epsilon} b_{\theta,j}(\beta) \right| \leq C_K (1 + \|\beta\|) \|s' - s\|. \quad (6.10)$$

Now, concerning the derivative of $\tilde{r}_j(\beta^j, b; \beta^{-j}, \theta)$, since

$$\log(\tilde{r}_j(\beta^j, b; \beta^{-j}, \theta)) = \frac{1}{2} \sum_{i=1}^n \left(\|y_i - K_p^{\tilde{\beta}_i} \alpha\|^2 - \|y_i - K_p^{\beta_i} \alpha\|^2 \right),$$

with $\tilde{\beta}_i = \beta_{i, b \rightarrow j}$, i corresponding to the i^{th} image, only one term of the previous sum is nonzero. We deduce from the fact that K_p is bounded and from (3.1) that $|\frac{d}{d\epsilon} \log(\tilde{r}_j(\beta^j, b; \beta^{-j}, \theta))| \leq C_K |\frac{d}{d\epsilon} \alpha| \leq C_K \|s - s'\|$, so that using the fact that $\tilde{r}_j(\beta^j, b; \beta^{-j}, \theta)$ is uniformly bounded for $\theta \in \hat{\theta}(\mathcal{K})$, $\beta \in \mathbb{R}^N$ and $b \in \mathbb{R}$, there exists a new constant C_K such that

$$\left| \frac{d}{d\epsilon} \tilde{r}_j(\beta^j, b; \beta^{-j}, \theta) \right| \leq C_K \|s - s'\|.$$

Thus, using (6.8), (6.9) and (6.10), we get for a new constant C_K that

$$\begin{aligned} & \left| \frac{d}{d\epsilon} r_j(\beta^j, b; \beta^{-j}, \theta) \exp \left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_\theta^2}{2} \right) \frac{\|e_j\|_\theta}{\sqrt{2\pi}} \right| \\ & \leq C_K (1 + \|\beta\|) \|s' - s\| (1 + |b - b_{\theta,j}(\beta)|)^2 \exp \left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_\theta^2}{2} \right) \frac{\|e_j\|_\theta}{\sqrt{2\pi}}. \end{aligned}$$

Since $\|f(\beta)\| \leq \|f\|_{V^p V^p(\beta)}$ and $V(a + b) = 1 + \|a + b\|^2 \leq 2(V(a) + V(b))$, we have $\|f(\beta_{0 \rightarrow j} + be_j)\| \leq C \|f\|_{V^p(V^p(\beta_{0 \rightarrow j}) + V^p(be_j))}$ with $C = 2^{2p-1}$. Hence, there exists an C_K such that $\forall (s, s') \in \mathcal{K}^2$, $\forall 1 \leq j \leq N$, $\forall \beta \in \mathbb{R}^N$ and $\forall \epsilon \in [0, 1]$:

$$\begin{aligned} & \int_{\mathbb{R}} \|f(\beta_{0 \rightarrow j} + be_j)\| \left| \frac{d}{d\epsilon} \left(r_j(\beta^j, b; \beta^{-j}, \theta) \exp \left(-\frac{(b - b_{\theta,j}(\beta))^2 \|e_j\|_\theta^2}{2} \right) \frac{\|e_j\|_\theta}{\sqrt{2\pi}} \right) \right| db \\ & \leq C_K \|f\|_{V^p V^p(\beta_{0 \rightarrow j})} (1 + \|\beta\|) \|s' - s\| \leq C_K \|f\|_{V^p V^p(\beta)} (1 + \|\beta\|) \|s' - s\|, \end{aligned}$$

where we have used the fact that a Gaussian variable has finite moments of all order. Since $(1 + \|\beta\|) \leq (2V(\beta))^{1/2}$, we get (updating C_K) that

$$\|(\Pi_{\hat{\theta}(1),j}^2 - \Pi_{\hat{\theta}(0),j}^2) f(\beta)\| \leq C_K \|f\|_{V^p V^{p+1/2}(\beta)} \|s' - s\|. \quad (6.11)$$

¹Note that the extra factor $(1 + \|\beta\|)$ appearing in the RHS of 6.10 compared to the RHS of 6.9 alleviate the need to show the usual Lipschitz condition $\|\Pi_{\hat{\theta}(s)} f - \Pi_{\hat{\theta}(s')} f\|_{V^{p'}} \leq C_K \|f\|_{V^q} \|s - s'\|$ with $q = p$. Weaker Lipschitz conditions as conditions **A3'** (ii) of Theorem 4.1 are needed

Now, looking at the first term in (6.3), we deduce easily from the previous study for $f \equiv f(\beta)$ that

$$\|(\Pi_{\theta(1),j}^1 - \Pi_{\theta(0),j}^1)f(\beta)\| \leq C_{\mathcal{K}}V(\beta)^{1/2}\|s' - s\|\|f(\beta)\| \leq C_{\mathcal{K}}\|f\|_{V^p}V^{p+1/2}(\beta)\|s' - s\|; \quad (6.12)$$

so that adding (6.11) and (6.12), we get (again updating $C_{\mathcal{K}}$) that

$$\|(\Pi_{\theta(1),j} - \Pi_{\theta(0),j})f\|_{V^{p+1/2}} \leq C_{\mathcal{K}}\|f\|_{V^p}\|s' - s\|. \quad (6.13)$$

We end the proof, saying that $\Pi_{\theta(1)} - \Pi_{\theta(0)} = \sum_{j=1}^N \Pi_{\theta(1),j+1,N} \circ (\Pi_{\theta(1),j} - \Pi_{\theta(0),j}) \circ \Pi_{\theta(0),1,j-1}$ where $\Pi_{\theta,q,r} = \Pi_{\theta,r} \circ \Pi_{\theta,r-1} \circ \dots \circ \Pi_{\theta,q}$ for any integer $q \leq r$ and any $\theta \in \Theta$ so that using (6.2) and (6.13), the result is straightforward. \square

LEMMA 6.5. *Let \mathcal{K} be a compact subset of \mathcal{S} . There exists a constant $C_{\mathcal{K}}$ such that for all $p \geq 1$ and any function $f \in \mathcal{L}_{V^p}$, $\forall (s, s') \in \mathcal{K}^2$, $\forall k \geq 0$, we have for $\theta = \hat{\theta}(s)$ and $\theta' = \hat{\theta}(s')$ that:*

$$\|\Pi_{\theta}^k f - \Pi_{\theta'}^k f\|_{V^{p+1/2}} \leq C_{\mathcal{K}}\|f\|_{V^p}\|s - s'\|.$$

Proof. We use the same decomposition of the difference as previously:

$$\Pi_{\theta}^k f - \Pi_{\theta'}^k f = \sum_{i=1}^{k-1} \Pi_{\theta}^i (\Pi_{\theta} - \Pi_{\theta'}) (\Pi_{\theta'}^{k-i-1} f - \pi_{\theta'}(f)).$$

Using Lemma 6.4, the fact that $\|\Pi_{\theta}^k(f - \pi_{\theta}(f))\|_{V^p} \leq \gamma_{\mathcal{K}}^k \|f\|_{V^p}$ with $\gamma_{\mathcal{K}} < 1$ (geometric ergodicity) and $\sup_{j \geq 0} \sup_{\theta \in \mathcal{K}} \|\Pi_{\theta}^j V^q\|_{V^q} < \infty$ we get:

$$\begin{aligned} \|\Pi_{\theta}^k f - \Pi_{\theta'}^k f\|_{V^{p+1/2}} &\leq C_{\mathcal{K}} \sum_{i=1}^{k-1} \|(\Pi_{\theta} - \Pi_{\theta'}) (\Pi_{\theta'}^{k-i-1} f - \pi_{\theta'}(f))\|_{V^{p+1/2}} \\ &\leq C_{\mathcal{K}} \|f\|_{V^p} |s - s'| \sum_{i=1}^{k-1} \gamma_{\mathcal{K}}^{k-i+1} \end{aligned}$$

and the lemma is proved. \square

We now prove that h is a Hölder function, adapting linearly Appendix B of [6].

Let $\beta \in \mathbb{R}^N$ and denote by $\theta = \hat{\theta}(s)$ and $\theta' = \hat{\theta}(s')$. Write $h(s) - h(s') = A(s, s') + B(s, s') + C(s, s')$, where

$$\begin{aligned} A(s, s') &= (h(s) - \Pi_{\theta}^k H_s(\beta)) + (\Pi_{\theta'}^k H_{s'}(\beta) - h(s')), \\ B(s, s') &= \Pi_{\theta}^k H_s(\beta) - \Pi_{\theta'}^k H_s(\beta), \\ C(s, s') &= \Pi_{\theta'}^k H_s(\beta) - \Pi_{\theta'}^k H_{s'}(\beta). \end{aligned}$$

Using the geometric ergodicity, Lemma 6.4 and Lemma 6.5, we get that there exists an $C > 0$, independent of k such that:

$$\begin{aligned} \|A(s, s')\| &\leq C \gamma_{\mathcal{K}}^k \sup_{S \in \mathcal{K}} \|H_S\|_V V(\beta), \\ \|B(s, s')\| &\leq C \sup_{S \in \mathcal{K}} \|H_S\|_V \|s - s'\| V^{3/2}(\beta), \\ \|C(s, s')\| &\leq C \sup_{S \in \mathcal{K}} \|H_S\|_V \|s - s'\| V(\beta). \end{aligned}$$

This yields

$$\|h(s) - h(s')\| \leq C V^{3/2}(\beta) (\gamma_{\mathcal{K}}^k + \|s - s'\|).$$

Hence, setting $k = \lceil \log \|s - s'\| / \log(\gamma) \rceil$ if $\|s - s'\| < 1$ and 1 otherwise, we get the result.

We can now end the proof of **(A3'(ii))**: On one hand we have:

$$\begin{aligned} \|(\Pi_{\theta}^k H_s(\beta) - h(s)) - (\Pi_{\theta'}^k H_{s'}(\beta) - h(s'))\| &\leq \|\Pi_{\theta}^k H_s(\beta) - \Pi_{\theta'}^k H_{s'}(\beta)\| \\ &\quad + \|\Pi_{\theta'}^k H_{s'}(\beta) - \Pi_{\theta}^k H_{s'}(\beta)\| + \|h(s) - h(s')\| \leq C\|s - s'\|V^{3/2}(\beta). \end{aligned}$$

On the other hand, we have thanks to the geometric ergodicity,

$$\|(\Pi_{\theta}^k H_s(\beta) - h(s)) - (\Pi_{\theta'}^k H_{s'}(\beta) - h(s'))\| \leq C\gamma^k V^{3/2}(\beta).$$

Hence for any $t \geq 0$ and $T \geq t$, we have

$$\begin{aligned} \|\Pi_{\theta}^t g_{\hat{\theta}(s)}(\beta) - \Pi_{\theta'}^t g_{\hat{\theta}(s')}(\beta)\| &\leq \sum_{k=t}^{\infty} \|(\Pi_{\theta}^k H_s(\beta) - h(s)) - (\Pi_{\theta'}^k H_{s'}(\beta) - h(s'))\| \leq \\ &\quad CV^{3/2}(\beta) \left[T\|s - s'\| + \frac{\gamma^{T+t}}{1-\gamma} \right]. \end{aligned}$$

Setting $T = \lceil \log \|s - s'\| / \log(\gamma) \rceil$ for $\|s - s'\| \leq \delta < 1$ and $T = t$ otherwise, using also the fact that for any $0 < a < 1$ we have $\|s - s'\| \log \|s - s'\| = o(\|s - s'\|^a)$, we get the result.

This proves condition **(A3'(ii))** for any $a < 1$.

We finally focus on the proof of **(A3'(iii))**. Once again we first prove a specific result for each function V_{θ} and obtain after a result for the function V .

LEMMA 6.6. *Let \mathcal{K} be a compact subset of \mathcal{S} and $p \geq 1$. There exists $C_{\mathcal{K},p} > 0$ such that for any $s, s' \in \mathcal{K}$, for any $\beta \in \mathbb{R}^N$,*

$$|V_{\hat{\theta}(s)}^p(\beta) - V_{\hat{\theta}(s')}^p(\beta)| \leq C_{\mathcal{K},p} \|s - s'\| V_{\hat{\theta}(s)}^p(\beta).$$

Proof. Indeed, there exists $C > 0$ such that for any $\hat{\theta}(s) = (\alpha, \Gamma_g)$ and $\hat{\theta}(s') = (\alpha', \Gamma'_g)$, $|\Gamma_g - \Gamma'_g| \leq C\|s - s'\|$. Therefore, there exists an C such that $\forall (s, s') \in \mathcal{K}^2$, $|\Gamma_g^{-1} - (\Gamma'_g)^{-1}| \leq C\|s - s'\|$ and

$$|V_{\hat{\theta}(s)}(\beta) - V_{\hat{\theta}(s')}(\beta)| \leq \sum_{i=1}^n \beta_i^t (\Gamma_g^{-1} - (\Gamma'_g)^{-1}) \beta_i \leq C\|s - s'\| V(\beta).$$

The result follows from the existence of a constant C such that $\frac{1}{c}V(\beta) \leq V_{\hat{\theta}(s)}(\beta) \leq CV(\beta)$ for any $(\beta, s) \in \mathbb{R}^N \times \mathcal{K}$. \square

LEMMA 6.7. *Let \mathcal{K} be a compact subset of \mathcal{S} and $p \geq 1$. There exist $\bar{\varepsilon} > 0$ and $C > 0$ such that for any sequence $\varepsilon = (\varepsilon_k)_{k \geq 0}$ such that $\varepsilon_k \leq \bar{\varepsilon}$ for k large enough, any sequence $\Delta = (\Delta_k)_{k \geq 0}$ and any $\beta \in \mathbb{R}^N$,*

$$\sup_{s \in \mathcal{K}} \sup_{k \geq 0} \mathbb{E}_{\beta,s}^{\Delta} [V^p(\beta_k) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k}] \leq CV^p(\beta).$$

Proof. Let K be a compact subset of Θ such that $\hat{\theta}(K) \subset K$. We note in the sequel, $\theta_k = \hat{\theta}(s_k)$. We have for $k \geq 2$, using the Markov property and Lemmas 6.2 and 6.6,

$$\begin{aligned} \mathbb{E}_{\beta,s}^{\Delta} [V_{\theta_{k-1}}^p(\beta_k) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k}] &\leq \mathbb{E}_{\beta,s}^{\Delta} [\Pi_{\theta_{k-1}} V_{\theta_{k-1}}^p(\beta_{k-1}) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k}] \\ &\leq \rho \left(\mathbb{E}_{\beta,s}^{\Delta} [V_{\theta_{k-2}}^p(\beta_{k-1}) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k}] + \mathbb{E}_{\beta,s}^{\Delta} [(V_{\theta_{k-1}}^p(\beta_{k-1}) - V_{\theta_{k-2}}^p(\beta_{k-1})) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k}] \right) + C \\ &\leq \rho \left(\mathbb{E}_{\beta,s}^{\Delta} [V_{\theta_{k-2}}^p(\beta_{k-1}) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k-1}] + C' \varepsilon_{k-1} \mathbb{E}_{\beta,s}^{\Delta} [V_{\theta_{k-2}}^p(\beta_{k-1}) \mathbf{1}_{\sigma(\mathcal{K}) \wedge \nu(\varepsilon) \geq k-1}] \right) + C. \end{aligned}$$

By induction, we show that

$$\mathbb{E}_{\beta,s}^{\Delta}[V_{\theta_{k-1}}^p(\beta_k)\mathbb{1}_{\sigma(\mathcal{K})\wedge\nu(\varepsilon)\geq k}] \leq \prod_{l=1}^{k-1}(\rho(1+C'\varepsilon_l))V_{\hat{\theta}(s)}^p(\beta) + \frac{C}{(1-\rho(1+C'\bar{\varepsilon}))}.$$

Choosing $\bar{\varepsilon}$ such that $\rho(1+C'\bar{\varepsilon}) < 1$ and introducing again $0 \leq c_1 \leq c_2$ such that $c_1V(\beta) \leq V_{\theta}(\beta) \leq c_2V(\beta)$ for any $(\beta, \theta) \in \mathbb{R}^N \times K$ end the proof. \square

This yields **(A3'(iii))**.

This concludes the demonstration of Theorem 4.4.

7. Conclusion and discussion. We have proposed a stochastic algorithm for constructing Bayesian non-rigid deformable models in the same context as [2] together with a proof of convergence toward a critical point of the observed likelihood. To the best of our best knowledge, this is the first theoretical result on convergence in the context of deformable template. The algorithm is based on a stochastic approximation of the EM algorithm using an MCMC approximation of the posterior distribution and truncation on random boundaries. Although our main contribution is theoretical, the preliminary experiments presented here on the US-postal database show that the stochastic approach can be easily implemented and is robust to noisy situations, yielding better results than the previous deterministic schemes.

Many interesting questions remain open. One may ask what is the convergence rate of such stochastic algorithms. A first result has been proved in [8] for the standard SAEM algorithm. Under mild conditions, the authors state a central limit theorem for an average sequence of the estimated parameters $(\theta_k)_k$. Concerning the generalization when introducing MCMC, a first step has been tackled in [5]. Under some restrictive assumptions the authors can prove a central limit theorem for an ergodic adaptive Monte Carlo Markov chain. We truly think that it is possible to obtain this kind of convergence rates for the SAEM-MCMC algorithm proposed in this paper.

Another question refers to the extension of the stochastic scheme to mixture of deformable models (defined as the multicomponent model in [2]) where the parameters are the weights of the individual components and for each component, the associated template and deformation law. This is of particular importance for real data analysis where the restriction to a unique deformable model could be too limiting. The design of such mixtures corresponds to some kind of deformation invariant clustering approach of the data which is a basic issue in any unsupervised data analysis scheme. This extension is, however, not as straightforward as it would appear at first glance: due to the high dimensional hidden deformation variables, a naive extension of the Markovian dynamics to the component variables will have extremely poor mixing properties leading to an impractical algorithm. A less straightforward extension involving multiple MCMC chains is under study.

Another interesting extension is to consider diffeomorphic mappings and not only displacement fields for the hidden deformation. This appears to be particularly interesting in the context of Computational Anatomy where a one to one correspondence between the template and the observation is usually needed and cannot be guaranteed with linear spline interpolation schemes. This extension could be done in principle using tangent models based on geodesic shooting in the spirit of [18].

REFERENCES

- [1] S. Allasonnière, Y. Amit, E. Kuhn, and A. Trouvé. Generative model and consistent estimation algorithms for non-rigid deformable models. In *IEEE Intern. Conf. on Acoustics, Speech, and Signal Processing*, volume 5, 2006.
- [2] S. Allasonnière, Y. Amit, and A. Trouvé. Toward a coherent statistical framework for dense deformable template estimation. *Journal of the Royal Statistical Society*, 69:3–29, 2007.
- [3] Y. Amit. Convergence properties of the Gibbs sampler for perturbations of Gaussians. *Ann. Statist.*, 24(1):122–140, 1996.
- [4] Y. Amit, U. Grenander, and M. Piccioni. Structural image restoration through deformable template. *Journal of the American Statistical Association*, 86(414):376–387, 1991.

- [5] C. Andrieu and É. Moulines. On the ergodicity properties of some adaptive MCMC algorithms. *Ann. Appl. Probab.*, 16(3):1462–1505, 2006.
- [6] C. Andrieu, É. Moulines, and P. Priouret. Stability of stochastic approximation under verifiable conditions. *SIAM J. Control Optim.*, 44(1):283–312 (electronic), 2005.
- [7] C. Chef d’Hotel, G. Hermosillo, and O. Faugeras. Variational methods for multimodal image matching. *International Journal of Computer Vision*, 50(3):329–343, 2002.
- [8] B. Delyon, M. Lavielle, and E. Moulines. Convergence of a stochastic approximation version of the EM algorithm. *Ann. Statist.*, 27(1):94–128, 1999.
- [9] A. P. Dempster, N. M. Laird, and D. B. Rubin. Maximum likelihood from incomplete data via the EM algorithm. *Journal of the Royal Statistical Society*, 1:1–22, 1977.
- [10] R. Douc, E. Moulines, and J. S. Rosenthal. Quantitative bounds on convergence of time-inhomogeneous Markov chains. *Ann. Appl. Probab.*, 14(4):1643–1665, 2004.
- [11] C. A. Glasbey and K. V. Mardia. A penalised likelihood approach to image warping. *Journal of the Royal Statistical Society, Series B*, 63:465–492, 2001.
- [12] U. Grenander and M. I. Miller. Computational anatomy: an emerging discipline. *Quarterly of Applied Mathematics*, LVI(4):617–694, 1998.
- [13] E. Kuhn and M. Lavielle. Coupling a stochastic approximation version of EM with an MCMC procedure. *ESAIM Probab. Stat.*, 8:115–131 (electronic), 2004.
- [14] S. Marsland, C. Twining, and C. Taylor. A minimum description length objective function for groupwise non rigid image registration. *Image and Vision Computing*, 2007.
- [15] S. P. Meyn and R. L. Tweedie. *Markov chains and stochastic stability*. Communications and Control Engineering Series. Springer-Verlag London Ltd., London, 1993.
- [16] F. Richard, A. Samson, and C. Cuénod. A saem algorithm for the estimation of template and deformation parameters in medical image sequences. *Statistics and Computing*, 2008.
- [17] C. Robert. *Méthodes de Monte Carlo par chaînes de Markov*. Statistique Mathématique et Probabilité. [Mathematical Statistics and Probability]. Éditions Économica, Paris, 1996.
- [18] M. Vaillant, I. Miller, M. A. Trouvé, and L. Younes. Statistics on diffeomorphisms via tangent space representations. *Neuroimage*, 23(S1):S161–S169, 2004.